


DISTRIBUTION STATEMENT A

Approved for Public Release
Distribution Unlimited

PHYSICAL ACOUSTICS SUMMER SCHOOL



ASTILOMAR CONFERENCE CENTER

VOLUME I: TRANSCRIPTS

DTIC QUALITY INSPECTED 2

1998 PHYSICAL ACOUSTICS SUMMER SCHOOL (PASS 98)

VOLUME 1: TRANSCRIPTS

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

This work relates to Department of Navy Grant 14-98-1-0044 issued by the Office of Naval Research. The United States Government has a royalty-free license throughout the world in all copyrightable material contained herein.

19991227 099

TABLE OF CONTENTS

FORWARD, LOGAN E. HARGROVE	1
LIST OF PARTICIPANTS	2
SCHEDULE	6
SONOLUMINESCENCE, ANTHONY A. ATCHLEY	7
SENSOR PHYSICS: SIGNALS AND NOISE, THOMAS B. GABRIELSON	38
THERMOACOUSTICS MADE SIMPLE, STEVEN L. GARRETT	73
NONLINEAR ACOUSTICS, MARK F. HAMILTON	111
CHAOS AND NONLINEAR BUBBLE DYNAMICS, WERNER LAUTERBORN	150
QUANTUM MECHANICS MINI-TUTORIAL, JULIAN D. MAYNARD	180
PERIOD, RANDOM AND QUASIPERIODIC MEDIA, JULIAN D. MAYNARD	203
RESONANT ULTRASOUND SPECTROSCOPY AND MATERIALS PHYSICS, ALBERT MIGLIORI.....	236
ACOUSTIC RESONATORS AND THE PROPERTIES OF GASES, MICHAEL MOLDOVER.....	270
SCANNING ACOUSTIC MICROSCOPY: LENSES, TIPS AND SONOELECTRONICS, CALVIN F. QUATE.....	302
POROUS MEDIA, JAMES M. SABATIER	322

1996 PHYSICAL ACOUSTICS SUMMER SCHOOL

FORWARD

These are the Proceedings of the 1998 Physical Acoustics Summer School (PASS 98). The lectures were recorded and most of the verbatim transcripts were subsequently edited by the authors for publication here.

Sponsored by the Office of Naval Research (ONR) and organized in cooperation with the Acoustical Society of America (ASA) and the National Center for Physical Acoustics (NCPA), PASS 98 was held June 14-21, 1998 at the Asilomar Conference Center in Pacific Grove, California, the site of the first and subsequent Summer Schools, PASS 92, PASS 94 and PASS 96. Participation in each of the Summer Schools was limited to a total of 50 that included students, lecturers, and discussion leaders.

The purpose of these Summer Schools is to bring graduate students, distinguished lecturers, and discussion leaders together to discuss a wide variety of subjects in physical acoustics. This gives the students the opportunity to meet experts and talk about topics most students ordinarily wouldn't encounter at their own colleges and universities. The focus was on graduate students and academic participants. Approximately half of the participants have been advanced graduate students in physical acoustics.

The Summer Schools have their beginning in an ONR Principal Investigators meeting in 1988 where it was decided that the best investment of this kind would be in a Summer School focused on graduate students. In 1990, an informal Summer School for students of ONR Principal Investigators and some invited guests was held as part of a Principal Investigator's meeting at the Naval Postgraduate School, Monterey, California, and this controlled experiment set the pattern and influenced the site selection for the subsequent PASS 92, PASS 94, PASS 96 and PASS 98. We hope that the Summer Schools will continue as biennial events with the high standards and wonderful success we have thus far enjoyed. This has happened and will happen again because everyone involved does their best, and for this I say, "Thank You!"

LOGAN E. HARGROVE – ONR 331

1998 PHYSICAL ACOUSTICS SUMMER SCHOOL

LIST OF PARTICIPANTS

-----LECTURERS-----

DR. ANTHONY A. ATCHLEY
PENNSYLVANIA STATE UNIVERSITY
GRADUATE PROGRAM IN ACOUSTICS
P.O. BOX 30
STATE COLLEGE, PA 16804
PHONE: 814-865-6364
FAX: 814-865-3119
EMAIL: AATCHLEY@PSU.EDU

DR. THOMAS B. GABRIELSON
PENNSYLVANIA STATE UNIVERSITY
APPLIED RESEARCH LABORATORY
0204M ARL BUILDING
STATE COLLEGE, PA 16802
PHONE: 814-865-1370
FAX: 814-863-8733
EMAIL: TBG@PSU.EDU

DR. STEVEN L. GARRETT
PENNSYLVANIA STATE UNIVERSITY
GRADUATE PROGRAM IN ACOUSTICS
P.O. BOX 30
STATE COLLEGE, PA 16804-0030
PHONE: 814-863-6373
FAX: 814-865-3119
EMAIL: GARRETT@SABINE.ACS.PSU.EDU

DR. MARK F. HAMILTON
UNIVERSITY OF TEXAS
DEPARTMENT OF MECHANICAL ENGINEERING
AUSTIN, TX 78712-1063
PHONE: 512-471-3055
EMAIL: HAMILTON@MAIL.UTEXAS.EDU

DR. ROBERT M. KEOLIAN
PENNSYLVANIA STATE UNIVERSITY
GRADUATE PROGRAM IN ACOUSTICS
P.O. BOX 30
STATE COLLEGE, PA 16804
PHONE: 814-865-1365
FAX: 814-863-8783
EMAIL: BONZO@SABINE.ACS.PSU.EDU

DR. WERNER H. LAUTERBORN
UNIVERSITÄT GÖTTINGEN
DRITTES PHYSIKALISCHES INSTITUT
BÜRGERSTRASSE 42-44
D-37073 GÖTTINGEN GERMANY
PHONE: 551-39-7713
FAX: 551-39-7720
EMAIL: LB@PHYSIK3.GWDG.DE

DR. JULIAN D. MAYNARD
PENNSYLVANIA STATE UNIVERSITY
DEPARTMENT OF PHYSICS
0330 DAVEY LABORATORY
UNIVERSITY PARK, PA 16802
PHONE: 814-865-6353
FAX: 814-865-3604
EMAIL: JDM@PSU.EDU

DR. ALBERT MIGLIORI
LOS ALAMOS NATIONAL LABORATORY
MS K764
P.O. BOX 1663
LOS ALAMOS, NM 87544
PHONE: 505-667-4133
FAX: 505-665-7652
EMAIL: MIGLIORI@LANL.GOV

DR. MICHAEL R. MOLDOVER
NATL. INST. OF STANDARDS & TECHNOLOGY
PHYSICAL AND CHEMICAL PROPERTIES DIV.
GAITHERSBURG, MD 20899-0001
PHONE: 301-975-2459
FAX: 301-869-4020
EMAIL: MICHAEL.MOLDOVER@NIST.GOV

DR. CALVIN F. QUATE
STANFORD UNIVERSITY
W.W. HANSEN LABORATORY OF PHYSICS
STANFORD, CA 94305-4085
PHONE: 650-723-0213
EMAIL: QUATE@EE.STANFORD.EDU

DR. JAMES M. SABATIER
UNIVERSITY OF MISSISSIPPI
NCPA
COLISEUM DRIVE
UNIVERSITY, MS 38677
PHONE: 662-915-5404
FAX: 662-915-7494
EMAIL: SABATIER@OLEMISS.EDU

----DISCUSSION LEADERS----

DR. HENRY E. BASS
UNIVERSITY OF MISSISSIPPI
NCPA
COLISEUM DRIVE
UNIVERSITY, MS 38677
PHONE: 662-915-5840
FAX: 662-915-7494
EMAIL: PABASS@OLEMISS.EDU

DR. CHRISTOPHER M. FORTUNKO
NATL. INST. OF STANDARDS & TECHNOLOGY
325 BROADWAY
BOULDER, CO 80303
PHONE: 303-497-3062
FAX: 303-497-5030
EMAIL: FORTUNKO@BOULDER.NIST.GOV

DR. LOGAN E. HARGROVE
OFFICE OF NAVAL RESEARCH
ONR 331 ROOM 503-13
800 NORTH QUINCY STREET
ARLINGTON, VA 22217-5660
PHONE: 703-696-4221
FAX: 703-696-6887
EMAIL: HARGROL@ONR.NAVY.MIL

DR. JERRY M. HARRIS
STANFORD UNIVERSITY
DEPARTMENT OF GEOPHYSICS
STANFORD, CA 94305
PHONE: 650-723-0496
EMAIL: HARRIS@PANGAEA.STANFORD.EDU

DR. CRAIG J. HICKEY
UNIVERSITY OF MISSISSIPPI
NCPA
COLISEUM DRIVE
UNIVERSITY, MS 38677
PHONE: 662-915-5963
FAX: 662-915-7494
EMAIL: CHICKEY@OLEMISS.EDU

DR. CRAIG W. HODGSON
LITTON GUIDANCE & CONTROL SYSTEMS
2436 VISTA WOOD CIRCLE, APT. #21
THOUSAND OAKS, CA 91362
PHONE: 818-715-4409
EMAIL: HODGSONC@LITTONGCS.COM

DR. THOMAS J. HOFER
NAVAL POSTGRADUATE SCHOOL
DEPARTMENT OF PHYSICS
CODE PH/HF
MONTEREY, CA 93943
PHONE: 408-656-2420
EMAIL: HOFER@PHYSICS.NPS.NAVY.MIL

DR. AMIT LAL
UNIVERSITY OF WISCONSIN
DEPT. OF ELECTRICAL & COMPUTER ENGR.
1415 ENGINEERING DRIVE
MADISON, WI 53706
PHONE: 608-265-9356
EMAIL: LAL@ENGR.WISC.EDU

DR. TOM J. MATULA
UNIVERSITY OF WASHINGTON
APPLIED PHYSICS LAB
1013 NE 40TH STREET
SEATTLE, WA 98105
PHONE: 206-685-7654
EMAIL: MATULA@APL.WASHINGTON.EDU

DR. RICHARD STERN
PENNSYLVANIA STATE UNIVERSITY
APPLIED RESEARCH LABORATORY
P.O. BOX 30
STATE COLLEGE, PA 16804
PHONE: 814-865-6344
EMAIL: RSTERN@PSU.EDU

DR. JOHN WENDOLOSKI
CSIRO TELECOMM. & INDUSTRIAL PHYSICS
BRADFIELD ROAD
WEST LINDFIELD
NSW 2070 AUSTRALIA
PHONE: 612-9413-7765
FAX: 612-9413-7200
EMAIL: JDRNCW@DAP.CSIRO.AU

-----STUDENTS-----

KEVIN BASTYR
PENNSYLVANIA STATE UNIVERSITY
314 WEST NITTANY AVENUE
STATE COLLEGE, PA 16801-4020
PHONE: 814-238-4223
EMAIL: KEVINBASTYR@PSU.EDU

FLORIAN BLONIGEN
WASHINGTON STATE UNIVERSITY
BOX 729 WEBSTER
PULLMAN, WA 99164-2814
PHONE: 509-335-1058

TODD BROOKS
YALE UNIVERSITY
DEPT. OF MECHANICAL ENGINEERING
P.O. BOX 208286
NEW HAVEN, CT 06520-8286
PHONE: 203-432-4362
EMAIL: TODD.BROOKS@YALE.EDU

NICHOLAS CERRUTI
WASHINGTON STATE UNIVERSITY
925 WEBSTER
DEPT. OF PHYSICS
PULLMAN, WA 99164-2814
PHONE: 509-335-7286
EMAIL: NCERRUTI@WSU.EDU

FANCESCO CURRA
UNIVERSITY OF WASHINGTON
1013 NE 40TH STREET
SEATTLE, WA 98105
PHONE: 206-543-9848
EMAIL: FCURRA@APL.WASHINGTON.EDU

BENJAMIN DZIKOWICZ
WASHINGTON STATE UNIVERSITY
DEPT. OF PHYSICS
PULLMAN, WA 99164
PHONE: 509-332-3557
EMAIL: DZIKOWIC@WSU.EDU

KEIR FOSTER
COLORADO STATE UNIVERSITY
905 WEST MOUNTAIN AVENUE
FORT COLLINS, CO 80521
PHONE: 970-493-3163
EMAIL: KFOZZI@LAMAR.COLOSTATE.EDU

JAMES HARDIMAN
UNIV. OF MASSACHUSETTS @ DARTMOUTH
P.O. Box 837
ASSONET, MA 02702-0837
PHONE: 508-644-9931
EMAIL: JHARDIMAN@UMASSD.EDU

STEPHEN D. HOLLAND
CORNELL UNIVERSITY
212 KIMBALL HALL
ITHACA, NY 14853
PHONE: 607-255-0986
EMAIL: SDH4@CORNELL.EDU

RONALD E. KUMON
UNIVERSITY OF TEXAS @ AUSTIN
1907 NUECES #46
AUSTIN, TX 78705
PHONE: 512-477-9704
EMAIL: KUMON@MAIL.UTEXAS.EDU

MICHAEL L. OELZE
UNIVERSITY OF MISSISSIPPI
NCPA
COLISEUM DRIVE
UNIVERSITY, MS 38677
PHONE: 662-915-5697
FAX: 662-915-7494
EMAIL: MLOELZE@OLEMISS.EDU

CLAUS-DIETER OHL
UNIVERSITÄT GÖTTINGEN
DRITTES PHYSIKALISCHES INSTITUT
BÜRGERSTRASSE 42-44
D-37073 GÖTTINGEN GERMANY
PHONE: 49-551-39-7723
EMAIL: OHL@PHYSIK3.GWDG.DE

TETSURO OISHI
UNIV. OF MASSACHUSETTS @ DARMOUTH
89 NEWTON STREET
NEWBEDFORD, MA 02740
PHONE: 508-984-0980
EMAIL: TOISHI@UMASSD.EDU

ANDI G. PETCULESCU
OHIO UNIVERSITY @ ATHENS
CLIPPINGER LABS
PHYSICS DEPT., ROOM 251
ATHENS, OHIO 45701
PHONE: 614-593-1710
EMAIL: APETCULE@HELIOS.PHY.OHIOU.EDU

GABRIELA PETCULESCU
OHIO UNIVERSITY @ ATHENS
CLIPPINGER LABS
PHYSICS DEPT., ROOM 251
ATHENS, OHIO 45701
PHONE: 614-593-1710
EMAIL: PETCULES@HELIOS.PHY.OHIOU.EDU

MATTHEW E. POESE
PENNSYLVANIA STATE UNIVERSITY
140 N. BARNARD ST., APT. #1
STATE COLLEGE, PA 16801
PHONE: 814-235-4891
EMAIL: MXP227@PSU.EDU

JENS K. POULSEN
AARHUS UNIVERSITET
INST. OF EXPERIMENTAL CLINICAL RESEARCH
SKEJBY SYGEHUS
DK-8200 ARHUS N DENMARK
PHONE: 45-89-49-6776
EMAIL: SKEJJKP@AAU.DK

WILLIAM SLATON
UNIVERSITY OF MISSISSIPPI
NCPA
COLISEUM DRIVE
UNIVERSITY, MS 38677
PHONE: 662-915-5635
FAX: 662-915-7494
EMAIL: WMSLATION@WAS.NET

GORDON SMITH
UNIVERSITY OF MISSISSIPPI
NCPA
COLISEUM DRIVE
UNIVERSITY, MS 38677
PHONE: 662-915-5635
FAX: 662-915-7494
EMAIL: SLIPSTK@OLEMISS.EDU

BRIAN D. STOREY
UNIVERSITY OF CALIFORNIA
6143 ETCHEVERRY HALL
BERKELEY, CA 94720
PHONE: 510-642-0204
EMAIL: BSTOREY@EULER.BERKELEY.EDU

KRIS VAN DE ROSTYNE
KATHOLIEKE UNIVERSITEIT LEUVEN
VITAL DE COSTERSTRAAT 68
B - 3000 LEUVEN
BELGIUM
PHONE: 75-93-59-44
EMAIL: KRIS.VANDEROSTYNE@FYS.KULEUVEN.AC.BE

RAY SCOTT WAKELAND
PENNSYLVANIA STATE UNIVERSITY
908 WEST BEAVER AVENUE, APT. B
STATE COLLEGE, PA 16801
PHONE: 814-861-3799
EMAIL: WAKELAND@PSU.EDU

BRIAN A. WATKINS
WASHINGTON STATE UNIVERSITY
P.O. Box 2263
PULLMAN, WA 99165-2263
PHONE: 509-334-7286
EMAIL: BWATKINS@WSU.EDU

MATTHEW WYCZALKOWSKI
UNIVERSITY OF CALIFORNIA
6143 ETCHEVERRY HALL
BERKELEY, CA 94720-1740
PHONE: 510-642-0204
EMAIL: MATWYCZ@ME.BERKELEY.EDU

STEVEN J. YOUNGHOUSE
UNIVERSITY OF TEXAS @ AUSTIN
4315 MARATHON BLVD.
AUSTIN, TX 78756
PHONE: 512-467-9357
EMAIL: SJY@MAIL.UTEXAS.EDU

YIBING ZHENG
YALE UNIVERSITY
DEPT. OF MECHANICAL ENGINEERING
P.O. Box 208286
NEW HAVEN, CT 06520-8286
PHONE: 203-432-4361
EMAIL: YIBING.ZHENG@YALE.EDU

WEI ZOU
CORNELL UNIVERSITY
212 KIMBALL HALL
ITHACA, NY 14853
PHONE: 607-255-0986
EMAIL: WEI@MSC.CORNELL.EDU

-----STAFF-----

LIBBY FURR
UNIVERSITY OF MISSISSIPPI
NCPA
COLISEUM DRIVE
UNIVERSITY, MS 38677
PHONE: 662-915-5808
FAX: 662-915-7494
EMAIL: LIBBY@OLEMISS.EDU

ROMELLE MILLION
MILLION REPORTING
P.O. Box 2419
ALEXANDRIA, VA 22301
PHONE: 703-836-3093

PASS 98

PHYSICAL ACOUSTICS SUMMER SCHOOL

ASLOMAR CONFERENCE CENTER • PACIFIC GROVE CA • 14-21 JUNE 1998

FINAL SCHEDULE

SUNDAY 14 JUNE	MONDAY 15 JUNE	TUESDAY 16 JUNE	WEDNESDAY 17 JUNE	THURSDAY 18 JUNE	FRIDAY 19 JUNE	SATURDAY 20 JUNE	SUNDAY 21 JUNE
	BREAKFAST 7:30 - 8:30	BREAKFAST 7:30 - 8:30	BREAKFAST 7:30 - 8:30	BREAKFAST 7:30 - 8:30	BREAKFAST 7:30 - 8:30	BREAKFAST 7:30 - 8:30	BREAKFAST 7:30 - 8:30
	A. A. ATCHLEY <i>Sonoluminescence</i> 8:30 - 11:30	FREE MORNING	M. F. HAMILTON <i>Nonlinear Acoustics</i> 8:30 - 11:30	C. F. QUATE <i>Scanning Acoustic Microscopy: Lenses, Tips & Sonoelectronics</i> 8:30 - 11:30	J. M. SABATIER <i>Porous Media</i> 8:30 - 11:30	M. R. MOLDOVER <i>Acoustic Resonators & the Properties of Gases</i> 8:30 - 11:30	CHECK OUT BEFORE NOON
	LUNCH 12:00 - 1:00	LUNCH 12:00 - 1:00	LUNCH 12:00 - 1:00	LUNCH 12:00 - 1:00	LUNCH 12:00 - 1:00	LUNCH 12:00 - 1:00	LUNCH 12:00 - 1:00
	T. G. GABRIELSON <i>Sensor Physics: Signals & Noise</i> 1:30 - 4:30	A. MIGLIORI <i>Resonant Ultrasound Spectroscopy & Materials Physics</i> 1:30 - 4:30	W. LAUTERBORN <i>Chaos & Nonlinear Bubble Dynamics</i> 1:30 - 4:30	FREE AFTERNOON	S. L. GARRETT <i>Thermoacoustics</i> 1:30 - 4:30	J. D. MAYNARD <i>Periodic, Random, & Quasiperiodic Media</i> 1:30 - 4:30	
REGISTRATION OPENS at 3:00							
DINNER 6:00 - 7:00	DINNER 6:00 - 7:00	DINNER 6:00 - 7:00	DINNER 6:00 - 7:00	DINNER 6:00 - 7:00	DINNER 6:00 - 7:00	DINNER 6:00 - 7:00	DINNER 6:00 - 7:00
INTRODUCTIONS J. D. MAYNARD <i>Quantum Mechanics Mini-tutorial</i> SOCIAL 7:30 - 10:30	First hour GABRIELSON <i>Sensor Electronics</i> DISCUSSION/SOCIAL 7:30 - 10:30	First hour MIGLIORI <i>Electronics for Acousticians</i> DISCUSSION/SOCIAL 7:30 - 10:30	INFORMAL DISCUSSION 7:30 - 8:30 BOONFIRE SOCIAL at PIT 9:00 - 11:00	R. M. KEOLIAN <i>Acoustics Demonstrations & SOCIAL</i> 7:30 - 10:30	INFORMAL DISCUSSION & SOCIAL 7:30 - 10:30	WRAP-UP, GRADUATION, & SOCIAL 7:30 - 10:30	

SPONSORED BY THE OFFICE OF NAVAL RESEARCH AND ORGANIZED IN COOPERATION WITH THE ACOUSTICAL SOCIETY OF AMERICA AND THE NATIONAL CENTER FOR PHYSICAL ACOUSTICS AT THE UNIVERSITY OF MISSISSIPPI

SONOLUMINESCENCE

Anthony A. Atchley
Graduate Program in Acoustics
Pennsylvania State University

[Transparency 1]

DR. ATCHLEY: First off, I would like to ask everyone in the audience who has worked or is now working in the field of sonoluminescence to stand up, please. [Members of the audience stood up.] All right, the rest of you look at these faces. When you want to ask questions, these are the people you want to ask.

I gave a similar talk at PASS two years ago and, when asked to do it, again, I thought it would be pretty easy. I still had all the transparencies. How much can change in two years, anyway?

A few months ago I decided to go to the library to catch up on a few things I thought I may have missed. I did a search of the Science Citation Index and found out that in the last 2 years -- almost starting from the time I gave the last talk -- there have been roughly 2 papers a month in reputable journals on sonoluminescence. A lot has changed.

If anyone has access to the transcripts from PASS'96, I began the lecture by saying that the students in the audience were, indeed, fortunate. At that time, they could spend an afternoon in the library and dig out the probably 15 or 20 important references in the field, read through them in a day or two, and be completely caught up in the area of sonoluminescence. You could follow this young field. You could understand pretty much all that was important in a very few number of days. Then, you could watch the literature as it progressed.

Well, I have to tell you that I do not think you students are fortunate as those of two years ago, because the field has really changed a lot. I also said at that time that you could put together an apparatus for sonoluminescence relatively inexpensively and do some important measurements.

In the last two years, the ante has been raised significantly. The sophistication of the equipment, the expense of the equipment, and the attention you have to pay to detail in the experimental aspects have really increased dramatically.

We know a lot more about sonoluminescence now than we did. You could easily paint a picture in which you think you know the answer, and I will try to do that, to some extent. But the field is changing rapidly. We can hope that in the next two years we will know what the answer really is. Then everyone who stood up will either be out of work or they will have found that there is a lot more wealth of information to come out of sonoluminescence. In two years, someone else will give a PASS lecture and tell us the important things that have happened since this lecture.

I should also point out that having Professor Lauterborn in the audience is a little bit daunting. This is equivalent, when talking about bubble dynamics, to trying to explain jet noise to Lighthill, or maybe milk to Lord Rayleigh -- I would rather not be doing it. (Speaking to Professor Lauterborn.) I apologize in advance for what I am going to do to your field.

This lecture is divided into 3 parts. You really cannot understand sonoluminescence until you understand how a bubble is trapped in a sound field, to begin with.

[Transparency 2]

I want to spend the first part of the lecture talking about acoustic levitation and bubble dynamics, to give you a feel for how you get a bubble to do this in the first place. The reason we talk about levitation and bubble dynamics is to get to a model of bubble dynamics that, if you just blindly apply, will give you some oversimplified but interesting predictions of what the pressure amplitudes and temperatures inside a collapsing bubble might be.

If you do not know anything about bubble dynamics and you apply this simple model, you can find a parameter regime where interesting things could happen. If you were going to anticipate something like sonoluminescence, you could focus on these regimes. That is the purpose of the first part of the lecture, to tell you about levitation, to tell you a little bit about how bubbles respond, and then what happens inside a bubble that makes things at all interesting.

Then I will talk about the early stages of what is called single-bubble sonoluminescence. Sonoluminescence as a field has been around for 60 years or so. It has been long known that if you insonify a liquid like water with high-amplitude sound waves, it will generate light.

It has been since only the early 1990s that the study of what is called single-bubble sonoluminescence really took off. I want to talk about the early days (meaning 5 or 6 years ago) of single-bubble sonoluminescence, just to lay the chronological map of how things evolved.

Then I want to talk a little bit about the parameter space that is important, and some experiments that people have performed to explore the parameter space of sonoluminescence. There is a wealth of parameters that can be tweaked and the parameter space is extremely large. I will talk a little bit about how people have probed that parameter space.

Then we will talk about theories. I have decided to approach the theories in the following way: I will talk about the essence of a theory in very simple terms (because I do not want to get into all the details), and then add more supportive experimental evidence to that. What I call theories is a mix of conjecture as to what is going on plus some experimental evidence. After this discussion we will be through.

I would like to invite the people in the audience who really do sonoluminescence to rebut or add to anything I have to say toward the end of the lecture.

Larry Crum wrote an article in *Physics Today* in September 1994. I bring this up partly because it is an interesting article, but there is a photograph of a former student of Sean Cordry. If you find that photograph, you will see Sean looking into a resonator where there should be a little bright spot, a little blue spot in the dark background.

Apparently, an editor decided the little bright spot in the dark background was a blemish in the photograph and retouched. So in the article, you see nothing cannot see it. What we think is interesting, other people think is a blemish in a photograph.

[Transparency 3]

A lot of the early work, after Felipe Gaitan did his initial work, was done by a group at UCLA. They have recently written an article in *Physics Reports* that I refer you to. It raises a lot of interesting questions and discusses some of the experiments they have performed. It is worth going through and is a great place to find references on sonoluminescence.

I should mention that as I go through my overheads and as you look in our notes, I have decided to make it almost a scrapbook of sonoluminescence literature. I tried to include the references to the papers and figures directly from the papers. Sometimes to get all that into one or two pages on a transparency I had to squash things, so things are sometimes too. However, you have the books in front of you.

My hope is that the notes serve will make it possible to find the articles to study later.

[Transparency 4]

This photograph shows a 500 mL spherical flask with piezoelectric ceramic cylinders glued to the sides. Those transducers are exciting this flask in one of the radial symmetric modes.

That sets up a radially symmetric standing wave inside this flask. If you inject a bubble into this system -- for instance, use a syringe to squirt some air in or maybe you have a piece of nichrome wire off to the side in the flask and you apply current to it so that it boils the fluid -- that bubble, or collection of bubbles, will be attracted toward the center of the flask, toward a pressure anti-node.

For the time being, forget that this bubble can give off light. The first thing you have to ask yourself is how can I trap a bubble in an acoustic standing wave? Why should that happen? This process is called acoustic levitation, and we will talk about that in a very simplified way here.

[Transparency 5]

Suppose there is no sound. I have a flask and I squirt a bubble in it. What is going to happen to that bubble?

PARTICIPANT: It is going to float to the top.

DR. ATCHLEY: Why? Because of buoyancy, right? So to keep that bubble in place (or it could dissolve, it depends on how big the bubble is) you have to counteract that buoyancy force.

The way you do that is with an acoustically induced force. The first question for homework is how do you calculate that acoustically induced force. You start with Newton's second law. If there is a buoyancy force, you have to counteract that with an acoustic force and you use Newton's second law to get that.

We are going to take a very simplified approach and assume there are only two forces that act on this bubble: the acoustic force and the buoyancy force. Let's see if we can figure out how to balance the buoyancy force with acoustics.

In addition, the other assumption is the following: Anything that is difficult about this problem I ignore. So we will do just the simplified stuff, to start with.

At equilibrium, you know that the time-averaged acoustic force has to equal the time-averaged buoyancy force. Let's calculate those two forces. The acoustic force is given by Eq. (2), which says that the force acting on a small volume, a small element, is going to be the volume of that element times the negative of the pressure gradient. That is the force equation, right?

The buoyancy force -- Eq. (3) -- is the density of the liquid times the volume of the bubble. Notice that both of the forces contain the volume of the bubble.

The pressure is given by Eq. (4), which says that there is some ambient pressure (the hydrostatic pressure in the fluid) and, superimposed on top of that, is the acoustic pressure. P_a is the pressure amplitude and I have assumed a standing wave form.

The volume of the bubble is written in a similar fashion in Eq. (5). It says that there is some mean volume (that is, the volume it has in the absence of sound) and an acoustically fluctuating part of the volume.

I have written it in such a way that V_a is the amplitude of the acoustically fluctuating component of the volume. It will be a function of the drive amplitude. The harder you drive a bubble (in simple terms), the more its volume is going to change. It is also going to depend on frequency. If you drive it near resonance V_a will be larger than if you drive it above or below resonance.

V_a is also going to depend on where the bubble is in the standing wave. For example, if you put it at a place where the pressure amplitude is 0, then it should not be compressed or expanded.

Now that we have the pressure and the volume, we can calculate the acoustic force and the buoyancy force. The acoustic force is given by Eq. (6), where I have substituted the expressions for the volume and the pressure.

(I should also point out, by the way, that when the transcript comes from these lectures, we get them back for editing. Here I can point to equations. It is obvious what I am referring to. But when I get the transcript, it will say "take this and put it 'here,' and substitute 'this,' and 'this is the result and then you do that.'" It shows how much you rely on visual clues during a lecture.

Calculating the acoustic force is straightforward. Next, you take the time average to get Eq. (7).

The buoyancy force is found from that Eq. (3). You take the volume of the bubble and substitute it. You take the time average and you find Eq. (9), which says that the time average buoyancy force is proportional to the average volume. Next, you equate the time-averaged acoustic force with the time-averaged buoyancy force and to get Eq. (11). Now, what does that say? I can solve Eq. (11) for z , the point where, for a given pressure amplitude and a given volume fluctuation of the bubble, the two forces balance. That will be the place where the bubble is trapped, the levitation point. It is found by equating the time-averaged acoustic pressure and the time-averaged buoyancy force.

The graph (which, in black and white, does not show up so well) has two graphs plotted. The z -axis is oriented vertically. For instance, this could represent a cylindrical or rectangular cross-section resonator oriented vertically, so increasing z is up.

The larger amplitude of the two curves is the acoustic amplitude as a function of position at, say, a particular instant in time. The smaller amplitude curve is this acoustic force. Suppose a bubble is placed at $z = 200$ arbitrary units. What would happen to it?

If you look at the acoustic force up on top axis, then at $z = 200$ the acoustic force is positive. If the acoustic force is positive, it will accelerate the bubble in the positive z direction.

It will keep going up and because it has inertia. It will overshoot the point where the acoustic force is zero. However, as it overshoots, the force becomes negative, which means the bubble will be drive back down. That point where the acoustic force is zero near $z = 250$ is a point of stable equilibrium. If the bubble is perturbed upward, the force is negative and pushes it back down. If the bubble is perturbed downward, the force is positive and pushes it back up. The bubble will be trapped at that point. Notice that that corresponds to a pressure anti-node.

That is acoustic levitation -- without gravity. If gravity is taken into account, the bubble is trapped above the antinode, according to Eq. (11).

Also notice that all the signs have worked out based on the assumption that the volume of the bubble goes down when the pressure goes up. That is the way you would normally think about things. If I have a bubble and I increase the pressure, the bubble ought to shrink. That is what forced me to choose that minus sign.

But you know from experience that if I have a harmonic oscillator and you ask for a given force what is the displacement response of the oscillator, the answer depends on whether the oscillator is driven above or below resonance.

A bubble is a harmonic oscillator. It has a resonance frequency and can be driven above or below resonance. This analysis is treating it as if we drive it below resonance; that is, when I increase the pressure, the bubble shrinks and when I decrease the pressure, the bubble expands. If I drive it above resonance frequency, then the opposite happens; the bubble will respond out of phase.

This is why the volume amplitude is a function of frequency. Bubbles driven above their resonance frequencies are trapped at the pressure nodes and not the pressure anti-nodes. Where the bubble is trapped depends on how it is driven.

MR. POESE: Does the success of trapping the bubble depend on the amplitude?

DR. ATCHLEY: Yes. If you look at the expression for the acoustic force, then it clearly is proportional to the pressure amplitude. If the pressure amplitude is small, then you cannot get enough force to counteract buoyancy.

If you look at Eq. (11). Suppose the pressure amplitude goes to 0. What happens to the right-hand side of the equation? It gets big. To solve Eq. (11) for z , the right-hand side of the expression has to be less than or equal to 1. For small drive amplitudes the right-hand side is bigger than 1 and there is no real value of z that satisfies the equation. That gives you a threshold for pressure amplitude.

DR. GARRETT: Another way to look at it is if you look at that corrugation up and down in force, if there is enough buoyancy force you can just ride up and down, whereas if the modulation is deep enough, you get stuck, you cannot make it over the hill.

In 0 G, everything there is amplitude-independent, because there are no other forces, but if there is some other force, it will tend to run it along this washboard. If it is bigger than the modulations in the washboard, then it is not going to travel.

MR. WAKELAND: How does the magnitude in VA compare -- can I say that VA has to be smaller in magnitude than --

DR. ATCHLEY: For linear theory it is small. For low-pressure amplitudes, then the fractional change in pressure is equal to the fractional change in volume, but as you drive the bubble harder and harder, its response is nonlinear -- we will see that in a couple of minutes.

MR. WAKELAND: But in general, acoustic levitation is not a nonlinear phenomenon?

DR. ATCHLEY: It is second order. You have to ask Mark Hamilton about nonlinear phenomena. Is this nonlinear or second order?

DR. HAMILTON: Like radiation pressure, it is energy, second order.

DR. ATCHLEY: It is not nonlinear, because everything I have done so far is linear acoustics. However, it does depend on the product of two acoustic quantities, $P-A$ and $V-A$, so it is second order in acoustic variables.

DR. LAUTERBORN: That is because it is proportional to P^2 .

DR. ATCHLEY: Right. It is a second-order phenomenon but it is not necessarily nonlinear. The theory I use is based on linear acoustics. It is like calculating energy in Kinsler and Frye, or intensity. That is purely linear acoustics, but it is a second-order phenomenon. Like thermoacoustics.

Any questions about levitation?

DR. MIGLIORI: Viscosity never comes into this? There is this problem of how you determine the onset of Rayleigh-Benard convection. It is a dynamics problem in which the viscous drag has to be matched to the energy delivered. As a novice, it is not obvious to me that you neglect that, but you say this works.

DR. ATCHLEY: This is enough to understand simplified acoustic levitation. There are certainly more complications, especially when you drive the bubbles more strongly. Professor Lauterborn will address all sorts of things, but if you are going to go into the lab and try to trap a bubble, this would get you a long way.

Now we have this bubble and it is trapped in this sound wave. What is that bubble going to do? Remember that not only is there an acoustic field that gives it a D.C. time-averaged force that traps it for hours on end, but the bubble is also experiencing an oscillating sound field. The bubble is going to respond with a fluctuating volume.

Let's try to understand how you might calculate, say, the radius of this bubble as a function of time.

[Transparency 6]

Let me put up this previous transparency and, Albert, read the title, "Simplified Acoustic Levitation." That means that all the difficult problems can be ignored. Now read the title of this one, "Simplified Bubble Dynamics." This title means that all the difficult problems will be deferred until Professor Lauterborn talks to us about bubble dynamics and chaos.

Let's understand a little bit about bubble dynamics, not only how you trap a bubble but how you drive it. The place you where begin is the Rayleigh-Plesset equation. Consider a bubble of instantaneous radius R , which is function of time.

This bubble is in an otherwise incompressible fluid. We do not worry about the fact that you actually have sound. We will neglect compressibility.

There are several approaches. One is to find the kinetic and potential energies and apply conservation of energy. The second was set up very nicely last night by Jay Maynard. You can also find the Lagrangian, which is the kinetic energy minus potential energy, and apply Lagrange's equation of motion there.

MR. WAKELAND: I was looking at this last night and eventually came to the following conclusion. I want you to tell me if this is true, that that is the kinetic energy and the potential energy of the whole system, which consists of the infinite fluid and the bubble.

DR. ATCHLEY: Absolutely right and we are about to see that.

DR. LAUTERBORN: I have a paper written in German with the Lagrangian formulas at home.

DR. ATCHLEY: You do? It works, right?

DR. LAUTERBORN: I can send it to you. It works.

DR. ATCHLEY: It brings up interesting details that get glossed over with this approach, I think. I would like to see that.

Let's talk about kinetic energy. Where is kinetic energy in this system? Decrease the pressure and let the bubble expand. Where is the kinetic energy? It is not in the bubble, not very much of it anyway. It is mainly in the fluid that is moving in response to the change in the bubble volume.

The kinetic energy is calculated by looking at the kinetic energy of shells of liquid surrounding the bubble. Find the kinetic energy density and integrate it over the whole fluid. Integrate the kinetic energy density from the radius of the bubble out to infinity. It is just a half times the density of the liquid times the velocity of one of these shells of liquid integrated over all volume.

To get the velocity of the fluid at some point, you have to relate u , the local velocity of the fluid to the velocity of the bubble wall. You do that by assuming the fluid is incompressible.

MR. WAKELAND: In reference to that last question, this kinetic energy is really from 0 to infinity and then you are saying that from 0 to R is small, because there is hardly any mass there, is that right?

DR. MIGLIORI: The other piece of it is that the surface of the container is so large that the far-field motion has no kinetic energy, if you integrate it out.

DR. ATCHLEY: Yes, because I am going to infinity.

DR. MIGLIORI: It is a little hydraulics problem, right? The surface is moving up and down slowly and then $1/2 mv^2$ is zero.

DR. ATCHLEY: Right. That is what Eq. (1) shows. If you do it in a spherical container, then you say that the mass flux for any shell of fluid is conserved. This says that the velocity goes as the ratio of the radii squared. As you get far enough away, the velocity drops off pretty quickly.

You substitute that expression for kinetic energy and now --

[Transparency 7]

-- you integrate. That is the result for u from the previous page put into the equation for kinetic energy.

What is the form of this equation? It is a half times something times the velocity squared. Kinetic energy is a half mv^2 , so this thing in the middle must be some mass-like term. Earlier in acoustics you probably learned (or maybe even in hydrodynamics) that the effective mass of a bubble driven below resonance is equal to 3 times the mass of the fluid displaced by the bubble. The effective mass of the bubble is 3 times the mass of the fluid displaced by the bubble.

Therefore, the kinetic energy can be written in the form, a half times some effective mass times the velocity of the surface squared. A simple result that could have been written down to start with, if you know what the effective mass of the bubble is.

Because energy is conserved, the time derivative of kinetic energy is equal to minus the time derivative of the potential energy. Take the time derivative, worry about the chain rule, and you get the result at the bottom of the page, where $\dot{\bullet}$ means time derivative, so \dot{R} is dr/dt (that is

the velocity of the bubble wall) and \ddot{R} is the acceleration of the bubble wall. Remember, R is the radius of the bubble.

[Transparency 8]

Potential energy is a little bit simpler. Suppose bubble is surrounded by some liquid. The pressure far from the bubble is P_∞ and the pressure at the bubble wall is P_L (the pressure in the liquid at the bubble wall). If these two pressures are not the same, then the bubble is going to change volume. The time rate of change of potential energy is the pressure difference times the change in volume dV/dt . The time derivative of potential energy is given by the equation at the bottom of the page.

[Transparency 9]

Now we can get to the Rayleigh-Plesset equation. The sum of the kinetic and potential energies is constant. Using the expressions we just found for the time derivatives leads to Eq. (1), Rayleigh-Plesset equation.

This is the first place you start on your way to understanding bubble dynamics. It is interesting that Rayleigh did this around 1917. So, about the time that the science that Jay talked about last night was going on, Rayleigh was thinking about bubble dynamics. He assumed that the pressure inside the bubble was 0. He had a collapsing void. Eventually, Milton Plesset came along and added more complexity.

Now, for Albert's edification, the details get pretty nasty as to how you calculate this P_L , the pressure just outside the bubble wall in the liquid. Clearly, there is going to be the pressure inside the bubble, which has a non-condensable gas component and a vapor component. Therefore, you have to worry about the equation of state of the gas inside the bubble to get the correct gas pressure as a function of radius.

One also has to worry about how the vapor pressure depends on what goes on inside the bubble. One also has to worry about surface tension, because if we go back to the first question I asked about what happens to a bubble when you put it inside a liquid, one answer is it dissolves. The reason is, because of surface tension the pressure inside the bubble is larger than the pressure in the liquid. As a result, the gas gets squeezed out of the bubble. In clean water, little bubbles do not last very long because of diffusion, so they may never have a chance to rise to the surface; they could just disappear.

One has to take viscosity into account, as well as other complicating factors.

MR. WAKELAND: So you have pressure of the gas and pressure of the vapor. The gas means like an argon bubble and the vapor means if it is in water, some amount of the water from the fluid that is in the bubble?

DR. ATCHLEY: Absolutely, water vapor, right.

MR. WAKELAND: And it can get out?

DR. ATCHLEY: It can condense on the bubble wall. The distinction is usually condensable gas and non-condensable gas.

DR. MIGLIORI: Does it collapse when the light comes off?

DR. ATCHLEY: When bubbles collapse, light comes off.

DR. MIGLIORI: No, no, is it true that essentially the volume goes almost to 0?

DR. ATCHLEY: Good question. It gets really small.

DR. MIGLIORI: Yes, but you know I want to pick on both viscosity and surface tension for small radii.

DR. ATCHLEY: Well, yes, but I am not going to let you pick on me about that. (Laughter)

You can pick on the experts. This is a nasty system. It is filled with complications and you are absolutely right, you could almost talk yourself into the state of being that anything of importance has been neglected. Nonetheless, interesting results come out.

So that explains what the liquid pressure is. The pressure at infinity is the ambient pressure plus some acoustic pressure. With modern computers and software packages like MatLab and Mathematica, you can just stick in these equations and pretty much believe the answer -- at least I can. I said "pretty much," I qualified myself, I am not a complete idiot.

It turns out that to really get very far in the field of bubble dynamics you cannot do much with the Rayleigh-Plesset equation. You have to take into account compressibility. After all, if you are driving this bubble with sound, the fluid has to be compressible. Compressibility effects make the derivation more complicated and I am not going to derive the Keller form of the Rayleigh-Plesset equation. I will say that a starting point for serious bubble calculations is Eq. (2), the Keller equation.

There are similarities between it and the Rayleigh-Plesset equation. For instance, you can see there is $R\ddot{R}$, $3/2 \dot{R}^2$. There are also liquid pressure terms. However, the biggest difference is that the compressibility of the fluid is included.

That introduces terms that are proportional to the Mach number of the bubble wall compared to the speed of sound in the fluid. Serious work in bubble dynamics starts with the Keller equation, or equivalent equations such as the Gilmore equation.

MR. POESE: I do not understand. In the Rayleigh-Plesset equation, you assume that the fluid is incompressible. How could P_L and P_{∞} be different?

DR. ATCHLEY: That is a good question.

MR. STOREY: I think the main reason you assume incompressibility is just so you can substitute velocity as a function of radius without having to solve full equations of motion. You just assume that knowing the radius of the bubble wall, you know the velocity anywhere.

DR. MIGLIORI: If I have an incompressible fluid in a beaker, the pressure at the top of the beaker and the pressure at the bottom of the beaker are different because of gravity.

DR. ATCHLEY: I think what Matt is worried about is if I turn sound on, since sound speed is infinite, how can P_L and P_{∞} be different?

The answer is that P_{∞} is the ambient pressure plus the acoustic pressure. I think it is true that if the fluid is incompressible, then in the absence of the bubble that would be same everywhere. The pressure on the wet side of the bubble is going to depend on what goes on inside the bubble.

DR. MIGLIORI: Is that used only to find the equilibrium volume of the bubble, basically?

DR. ATCHLEY: No, not equilibrium. It is the dynamic volume of the bubble. R is a function of time.

DR. MIGLIORI: No, I meant the incompressible piece.

DR. ATCHLEY: The incompressible piece allows me to fit this derivation onto 3 pages. (Laughter)

The purpose of assuming incompressibility is to allow me to use the expression for u as a function of r at the bottom of *Transparency 6*.

DR. LAUTERBORN: I think the main answer is that the pressure immediately is applied to the surface of the bubble. If you have compressibility, there is a time lag. If you apply a sound field it needs some time to reach the bubble surface. That is the main difference. Incompressibility means you instantly apply it to the bubble surface without time delay.

DR. MIGLIORI: I will need to look at that later. I do not quite understand it.

DR. LAUTERBORN: It need not be in equilibrium. It is only in dynamic equilibrium.

DR. ATCHLEY: Yes, it is equilibrium enough to simulate what goes on inside the bubble.

MR. WAKELAND: So P_1 became P_0 plus P_{acoustic} and then, in the Keller equation, the P_{acoustic} became P_s , is that right?

DR. ATCHLEY: Yes.

MR. WAKELAND: And c is the speed of sound in the liquid and, therefore, that is the constant that includes the compressibility of the liquid, is that right?

DR. ATCHLEY: Well, it does, but other terms have compressibility effects, too.

MR. WAKELAND: c is the speed of sound in the liquid?

DR. ATCHLEY: Yes.

DR. LAUTERBORN: At the surface of the bubble.

MR. WAKELAND: At the surface? It is not the same everywhere.

DR. ATCHLEY: No.

MR. WAKELAND: We are not doing linear acoustics?

DR. ATCHLEY: It gets nasty quick.

DR. LAUTERBORN: It also has a time lag in. It is very complicated.

DR. ATCHLEY: Right. The point I want to make is that you can see the origins of these equations back in Rayleigh-Plesset and the fundamentals of that are kinetic and potential energies. If you are careful about calculating those things properly, you can get more complicated results, but we do not have time to do that.

[*Transparency 10*]

Suppose you take the Keller equation, assume a bubble of a certain size (I assumed this was 5 μm , which is a pretty tiny bubble, by the way), and drive it at 3 different pressure amplitudes.

The upper graph shows the acoustic pressure as a function of time, normalized so that it peaks out at 1. It is on this graph to show you the phasing of the pressure-to-bubble response.

The next 3 graphs show the response of the bubble to 3 different pressure amplitudes; 0.05 atm, so it is 5% of an atmosphere, 50 %, and 1.25 times the ambient pressure in the fluid, assumed to be 1 atm. The drive frequency is 25 kHz.

Consider the small drive amplitude case. Notice that if you compare this graph to the pressure, it behaves linearly to a good approximation. That is, when the pressure goes up, the bubble gets smaller and when the pressure goes down, the bubble gets bigger. If you look at the change in radius, it goes up and down about half a micron. That is in proportion to the pressure amplitude. If you look at the frequency spectrum, it looks pretty linear. There are no harmonics; it is just a bubble getting larger and smaller.

If you drive the bubble harder at half an atmosphere, then you see that the response is not linear at all. The expansion part of the bubble cycle looks a lot different from the compression

part of the bubble cycle. It gets bigger than it does smaller and it is not a sinusoidal function any more.

Then if you drive it really hard at 1.25 atm, the response is very nonlinear.

DR. LAUTERBORN: It also gets asymmetric in compression and expansion. It is more expanded than compressed.

DR. ATCHLEY: Right, it gets bigger than it does smaller.

The bubble starts off at 5 μm , gets much bigger, maybe by a factor of, in this graph, 7 times radius, so 7-cubed in volume. Then it collapses to a very small size and rebounds, bounces up and down, stays near the equilibrium radius of 5 μm , then undergoes this highly nonlinear oscillation, again.

DR. GABRIELSON: If you used the non-Keller form of the Rayleigh-Plesset equation to do these, would there be significant differences in form or is it just a matter of numbers?

DR. ATCHLEY: I did this a couple of years ago now. It seems to me that I had trouble with the stability of the solution using Rayleigh-Plesset. Frankly, I do not remember.

MR. STOREY: The main difference comes with collapse. Sometimes the equations will just blow up and you cannot numerically solve them any more. The expansion will be exact the same because the Mach number is very small.

DR. ATCHLEY: Yes that is a good point. The initial expansion is relatively slow in the sense that the Mach number is small. Therefore, Rayleigh-Plesset ought to be pretty good. During the final stages of the collapse, the validity falls apart.

MR. SMITH: Is that ringing a function of the bubble's natural oscillation?

DR. ATCHLEY: Yes.

MR. STOREY: But the decay of that ringing has nothing to do with dissipation, right?

DR. ATCHLEY: Viscosity is included. So is surface tension and radiation, the dp_L/dt term.

MS. ZOU: It looks like the acoustic period is long enough so that the ringing dies out and the radius becomes steady. What if your period is really small?

DR. ATCHLEY: The question is, if you look at the graph, before the bubble starts to expand again, these oscillations have decayed out. The question is what happens if the acoustic period were shorter so that they would not die out?

MS. ZOU: Right.

DR. ATCHLEY: That is a good question for Professor Lauterborn. (Laughter)

DR. LAUTERBORN: I will show some graphs on how we approach that to consider the whole frequency range. I will give a survey on that.

DR. ATCHLEY: I should point out that the time axis here starts at 60 μsec . The period of 25 kHz acoustic signal is 40- μsec . I have waited a cycle, or maybe more than a cycle, to plot this. Transients appear in the solution that have to decay out. These results are close to steady state.

If you change the period, it might not be too wrong to say that the collapse will happen pretty much the same and you will get some rebounds. However, now the condition will be different when you start to drive it back up again, so the response should look different.

DR. LAUTERBORN: But then you are not far below the resonance, because these oscillations are the resonance oscillations of the bubble at that pressure.

DR. ATCHLEY: The picture now is that if you use the Keller equation, you can get a wealth of behavior from a simple system. This is just a bubble in a fluid being driven with a sound field. The type of response you get is extremely nonlinear. Recall that this graph is the radius of the bubble wall as a function of time.

What happens inside the bubble? Things can get really complicated. It is simple to do if you assume a simple equation of state. For instance, assume an ideal gas, or a van der Waals gas, and some nice thermodynamic process, adiabatic or isothermal. If you do not worry too much about what happens to the vapor inside the bubble and you do not worry about nasty things, such as viscosity or surface tension as things get small, if you just blindly put in an equation of state and do not worry about problems, you can look inside the bubble.

I need to go back and remind everybody where we are on the outline. We are at this point: oversimplified unbelievable predictions. The reason to do that is just to set the stage. The next figures I am about to show should not have too much faith placed in them. Their purpose is to point us to an interesting parameter region.

[Transparency 11]

The graph in the upper left is the last figure I showed you, that is, a 5-um bubble being driven at 1.25 atm of pressure. This is a result of the Keller equation. You can take the time derivative and get the velocity of the bubble wall (upper right). You can assume ideal gas, maybe adiabatic behavior, inside the bubble and get a pressure inside the bubble (lower left). You can also get a temperature inside the bubble (lower right).

The bubble expands, has a large expansion ratio, and then collapses rapidly, quickly, violently. What you see is that during this collapse the wall velocity gets extremely large. Remember, the Keller equation is valid to first order in Mach number. The speed of sound in fluid is 1500 m/sec for water, and you can see these velocities are way too big for this equation to be valid. The point is that the bubble wall is moving pretty quickly during the final stages of the collapse.

The internal pressure spikes when the bubble collapses spikes. You cannot believe the numbers on the graph, but they indicate that during the collapse the contents of the bubble get compressed a lot.

If you take just simple adiabatic heating, the temperature also goes up a lot. I ask you, if you went back to your lab, built a levitator, put a bubble in there, drove it, and got these highly nonlinear oscillations, where would you look in the acoustic cycle for something interesting to happen?

Well, maybe during the collapse, where simple analysis tells you that the pressures and temperatures get extremely large. Therefore, this points us to a place to look for something interesting to happen, and now I can really talk about sonoluminescence.

MR. ZHENG: Can anybody really measure the temperature in the bubble?

DR. ATCHLEY: No.

MR. ZHENG: I think it is very hot.

DR. ATCHLEY: I agree. No one knows, really, what is inside the bubble.

I think that at this point we should take a break for about 10 minutes. After the break, we can talk about sonoluminescence. Finally, after an hour of bubble dynamics, the stage has been set. We know where we are going to look for something interesting to happen. Let's come back, look there, and see what happens.

DR. ATCHLEY: Part two. I would like to encourage you to keep asking questions. I have a lot of material in the notes, but we do not have to cover it all. I can pretty much stop when the time runs out, so do not worry about asking questions. I would rather get the audience involved.

There is another system that you can perturb with an external field and the response depends on what is inside the system, and that is called a volleyball. You have this thing that has gas inside and you hit it with an impulse and it moves. Jay Maynard wanted me to remind you that you can do experiments with volleyballs at 11:30 on the volleyball court. Jay will tell you where that is.

[return to Transparency 5]

Let me address some questions that came up during the break. Remember that acoustic levitation depends on the acoustic component of the volume of the bubble.

[return to Transparency 10]

Now, recall what we did just before the break. The volume is a very complicated function of time. To really understand what is happening when trying to trap and excite a bubble, that you must take the actual response into account when calculating that trapping point. That is the subject of some nice papers, one by Tom Matula and his collaborators, and also by Professor Lauterborn and his colleagues.

[Transparency 12]

I will remind you that although sonoluminescence has been known for decades, things got particularly interesting in the early 1990s when Felipe Gaitan was a student at the University of Mississippi. It is coincidental that the first paper that I know of on single-bubble sonoluminescence was presented at an International Symposium on Nonlinear Acoustics in Austin, Texas, organized by Mark Hamilton and David Blackstock in 1990. Professor Lauterborn is organizing another ISNA in Göttingen. There will be something just as startling as sonoluminescence that comes out of that ISNA.

[Transparency 13]

Felipe took a bubble, trapped it in a resonator, and used laser scattering from the bubble to get a measure of the radius of the bubble as a function of time. He saw this interesting result.

[Transparency 14]

This upper curve shows the pressure amplitude in the acoustic resonator as a function of time. The middle curve shows what at that time was a pretty good, but by today's standards a pretty crude, representation of the bubble response f as determined by laser scattering. You can see that the general features of what we just calculated are there. The bubble gets very large and collapses, gets very large, and collapses, etc.

Felipe found that, if you look at the bubble with a photomultiplier tube, every time the bubble collapses, the photodetector produced a spike -- one spike every time the bubble collapses. This behavior repeats itself time after time after time. This is single-bubble sonoluminescence. It occurs in the parameter range that we thought might be interesting during the first part of the lecture.

[Transparency 15]

Felipe went on to make a number of fundamental measurements. For instance, he showed that there is only a limited range of drive amplitudes for which you can get this phenomenon. I need to take a few minutes to explain how Felipe got bubbles to luminesce in his original experiments.

There are a number of variations on the theme, but I think it is important to tell you how the bubble evolves. You take your resonator filled with clean water and you excite a standing wave in that resonator. You inject some air into that resonator. Of course, when you first squirt air into it, you get a lot of bubbles that rise right to the top.

If you have the sound field turned up, many bubbles will be attracted toward the pressure anti-node and coalesce. At low amplitudes, you have a pretty big bubble, stably levitated at the pressure anti-node, and it will stay there.

Remember that there are really two forces going on. There is the time-averaged acoustic force, which traps the bubble, and there is the instantaneous force exerted by the drive pressure. The levitating field is also making the bubble respond. That region is the lower part of the figure. You get low-amplitude, radial pulsations of the bubble similar to the second graph on Transparency 10.

Now turn up the amplitude. The first thing you notice is that the bubble just oscillates with larger amplitude. If you measure the radius of the bubble as a function of time, then you will find that the amplitude will get a little bit distorted, but it is still pretty well behaved.

I should also tell you, by the way, that Felipe used a degassed, mixture of glycerin and water (that was the first fluid he used). You turn the amplitude up more and more and the bubble behavior gets really crazy; it starts moving around in the sound field and you think it is just going to rip apart.

You crank it up a little bit more in amplitude and the behavior just gets wild. It dances all over the place. It looks like it is shedding little bitty bubbles and they are being shot around in the fluid. This is the region called non-spherical pulsations. The bubble is going crazy. Turn the amplitude up more and the bubble stabilizes rock solid, gets really small, barely visible to the eye, and it is glows.

When I was first doing sonoluminescence that is how I could tell it was glowing. Because I had the lights on in the room I couldn't see the luminescence. I would inject the bubble and take it through this evolutionary process. The bubble would stabilize, rock solid, and appear smaller to the eye. If I turned out the lights, the bubble would be glowing. Eventually, I could see that it was glowing even with the room lights on. That puts us at the lower stability threshold.

Then, as you turn the amplitude up more and more, the light gets brighter and brighter, visible to the naked eye. When I was doing sonoluminescence at the Postgraduate School, often I would have admirals come around for VIP tours. I was always hesitant to show them sonoluminescence, because as you get older, your visual acuity starts to drop. So some admiral with stars and medals everywhere comes in. I would tell them what going on and point to this glowing bubble. Sometimes the admiral could not see it. (Laughter)

Anyway, some people can see it with the naked eye, for others it is a little bit tougher.

You turn the amplitude up more and more, the bubble generally gets brighter and brighter. Then you turn the amplitude up further and the bubble disappears. It gets shot out of the sound field. The bubble goes away and the light, of course, quits.

MS. ZOU: I look at the radius of this bubble and it is like 15 mm. Is that reasonable, that big of a bubble?

DR. ATCHLEY: That has got to be micron. I never noticed that. It has got to be microns, yes. There is a typo in your volume, Mark. If you had a 15-mm bubble doing this, you could maybe cook a chicken with it. (Laughter)

Felipe found that there was only a certain narrow parameter range where a bubble would luminesce. He made a lot of other measurements, too.

About that time Seth Putterman visited NCPA and was extremely fascinated by the process. He is a much better physicist than I or many people in the room and he immediately saw that this was a really curious phenomenon.

How could a relatively simple classical system take acoustic energy, which, mechanically, is very weak and concentrate it enough to emit light?

Seth had a very talented graduate student at the time, Brad Barber. Felipe graciously showed them how to produce sonoluminescence and they went back and did some really first-class work.

I am going to spend the next few minutes talking largely about the UCLA work, because many of the papers in the early 1990s were from UCLA.

[Transparency 16]

This graph is from a paper by Brad Barber and Seth Putterman published in Nature, 7 years ago now. They found that, just as Felipe showed, there is only a limited range of pressure amplitudes over which bubbles luminesce. They measured the threshold as a function of drive frequency. I should remind you that because you are operating in a resonator, you cannot continuously change frequency over any value you want. You are restricted to the resonant modes of the cavity. That is why these frequencies are discrete.

They also showed that the amount of light emitted per burst was something like 100,000 or 1,000,000 photons, which is a pretty good count (although, maybe, I cannot say you can count photons based on Jay's lecture last night). If you use a detector in the way they used it, when light hits that detector, it sends a blip -- not that light is a particle -- the detector gives a pulse. If and you count those pulses the number is about 100,000 or 1,000,000.

[Transparency 17]

They also found that the emission was very regular. The distribution of pulse heights from the luminescence can be very stable. The amplitude is the same just about all the time.

[Transparency 18]

They also found that the jitter from pulse to pulse is also small. This graph shows the arrival time between successive flashes. These are really small time differences. That multiplier is 10^{-6} , I believe, and there are a lot of digits in front of that. The scale is 200 psec/division. This shows that the jitter from pulse to pulse is very small, about 50 psec, which is extremely short.

That is pretty nice. You can get a classical system emitting light very synchronously, very stably. However, the field really got -- I do not want to say hot -- really became more active when they asked the question, how short is the pulse?

[Transparency 19]

This graph shows four panels. The upper left panel shows the response of a photomultiplier tube that has a 2-nsec rise time to sonoluminescence. The lower left graph is the output from the same photomultiplier tube exposed to a 34-psec laser pulse. The two responses are the same, which means they were measuring the response of the tube and not the rise time of the pulse.

They got a faster photomultiplier tube, about a factor of 10 faster, a rise time of about 250 psec and did the same experiment. They get more or less the same result. They were measuring the response of the tube and not the duration of the flash.

[Transparency 20]

So what do you do? You get a faster photomultiplier tube. Except they don't come much faster, maybe only a factor or 2 faster, from about 240 psec to about 170 psec rise time. They repeated the same experiment and found that the rise time of the response of the tube to sonoluminescence and to the laser pulse was essentially the same.

The fall time difference is due to the fact that the laser does not turn off cleanly. They found very little difference in the rise time. They inferred after a lot of signal processing that there was a 4-psec difference in the rise times. From that they inferred a limit to the pulse duration of about 50 psec, which is extremely short. The picture is the following.

[Transparency 21]

This is a map of the California coast. This little part here is my former place of employment, the Naval Postgraduate School -- we are somewhere over here. Suppose you went back to where I used to work and put a sonoluminescence device on top of the main building and look at it from a satellite in space. You would see spherical waves emanating, flashes that radiate out symmetrically in all directions.

The spacing between the pulses, if you drive it at 25 kHz, is about 15 km. The width of the lines is about as wide as your thumb, because a 50 psec pulse of light is not very long. The precision with which the lines are spaced is also about the width of your thumb.

This graph indicates that sonoluminescence extremely stable. The pulse spacing is long in time but the pulse length is very small because the pulse is so short. The jitter is also extremely small. That is a pretty nice thing to study.

How can you take this classical acoustic system and with state-of-the-art photodetectors not be able to resolve the pulse duration? That is worth studying.

MR. WAKELAND: Is jitter from the driver an issue at all in making this measurement? Can we say that the driver is free of jitter?

DR. ATCHLEY: No. In fact, if you look at the specs of the H.P. function generator, the jitter is a lot bigger than this. However, the system itself has a relatively high Q, so you have a high-Q system driven with something with a little bit of jitter, does that explain it all? Maybe.

[Transparency 22]

We talked about the duration of the flash. If you look at it with your eye, it is bluish; the question is what is the optical spectrum? Bob Hiller measured the spectrum of the light and it turns out that sonoluminescence is extremely frustrating for the following reasons.

It is bright enough to see with your eye, but it is dim enough so that a lot of really fast detectors will not detect it. It is also out of the spectral range where a lot of really fast photodiodes work. It also happens inside a liquid. This graph shows the wavelength of light and a measure of the intensity of the flash (this is optical spectrum). The visible range is somewhere over on the right side of the graph. Most of the radiation is in the ultraviolet.

If you look at the data, they drop off 200-240 nm in the ultraviolet. Water starts to absorb light at around 200 nm or so. Also ordinary optics do not work as well in the ultraviolet. What you think of as a mirror in the visible range is not a mirror in the ultraviolet range and what you think of as a lens in the visible range is not a lens in the ultraviolet.

You can get sophisticated and push the threshold down to maybe around 200 nm, but you cannot do a lot more with water and glass. It is really frustrating. It is bright enough to see with your eye, but too dim for fast detectors to work. Further, the interesting part of the spectrum appears to be in a range that is largely inaccessible. The question is how much more light is there in the ultraviolet.

DR. SABATIER: Has anybody been able to push something like a metal-bearing compound in the bubble and thereby see line spectra. Does that give you a line spectrum from which you can --?

DR. ATCHLEY: We will talk about the spectra. It turns out that for the old-fashioned sonoluminescence from multiple bubbles, one of the characteristics of the spectrum (we will see this later) is that you see lines that depend on what you put in the water. If you do it in salt water, you see sodium lines. If you do it some oil or some chemical, you will see carbon lines, all sorts of things. If you do single-bubble sonoluminescence in the same liquids, you do not see lines.

The dashed line is the raw data. The dots are the data that have been corrected for as many things as they can correct for. The solid line is a blackbody fit. In the early days of sonoluminescence people would obtain a spectrum and say, well, let's try to fit it to something. What can we fit it to? A blackbody. Therefore, they did. There is no reason to believe that it is a blackbody, necessarily, but you can fit to it. The tail fits pretty well. If you look at the temperature that the fit corresponds to, it is something like 25,000 degrees, which is also interesting.

People who do spectroscopy say that it is not a blackbody unless you see a tail and a peak and a decay. You cannot tell from just the tail what the animal is. Nonetheless, people fit to it and they got extremely high temperatures, like 25,000-30,000 degrees.

Now, you have a little bubble, a nice simple system, inexpensive to operate, generates extremely short pulses of light, and a lot of energy is in the ultraviolet.

DR. MAYNARD: At 25,000 K, where is the peak in the blackbody?

DR. ATCHLEY: Oh, I forget. An eV is about 12,000 degrees, so 25,000 must be a couple of eV, which is somewhere in the U.V., maybe.

DR. MIGLIORI: Well, eV is visible light, so I do not know where the peak is, but it is not going to be too far off of that. It is probably very close to where your data stopped.

When you are fitting the tail of blackbody radiation, the absolute intensity is crucial or else you have diluted blackbody radiation. You do not really quite know what is going on. Obviously, you have problems with getting the light out, reflective surfaces, and so on. In addition, if you are looking at only the tail, it is very hard to fit those accurately. Do you know what the error bars are?

DR. ATCHLEY: There are two parameters to fit. One is the temperature and one is the size. They did not report anything about the size. We measured spectra and tried to play around with the size of the source. We did not get anywhere.

DR. MIGLIORI: But if you tried to fit a 50,000 K blackbody spectrum to that tail, how much worse would it look, for example?

DR. ATCHLEY: When we were playing around with it, 5000 degrees made a difference. It looks different enough that people would believe that maybe 25,000, if it were blackbody, would be right. We got 20,000 for our spectra; maybe you would get 30,000. It is all in that ballpark.

DR. MIGLIORI: But not 50,000 or 10,000?

DR. ATCHLEY: No. There was enough difference there that you could tell.

DR. FORTUNKO: Does temperature even have a meaning in this context?

DR. ATCHLEY: I should throw out some more caveats. I am going to try to be very noncommittal on hypotheses of what is going on. What I would like to do is present mainly experimental data to show what people think is true. What goes on inside the bubble is extremely complicated and all these questions come up: does temperature mean anything, things are happening this quickly, can you even use equilibrium thermodynamics in the first place? I want to avoid all these complications because nobody knows yet.

DR. MIGLIORI: The speed of light is 0.3 mm/psec, so say it is a 50-psec flash, you could make an argument that it is in full thermal equilibrium, because you had many transit times of photons inside that bubble during the flash of the light, so maybe temperature is not so bad.

DR. ATCHLEY: It also depends on the electron collision times, too. In serious calculations of what is going on inside a bubble they try to take into account the photon temperature, the electron temperature, and the molecular temperature.

MR. STOREY: Yes, you can kind of show by order of magnitude that the mean free path is so small that the thermodynamic equilibrium assumption is okay.

DR. ATCHLEY: Yes.

[Transparency 23]

What happens if you cool the water? The data I showed previously was for room temperature water. They cooled the water down to, say, 10 degrees C, a change of 10 on an absolute scale, a pretty small change in temperature, and found that the spectrum for 10-degree water was different from 20-degree water. The bubble was brighter and the spectrum shifted more toward the UV.

[Transparency 24]

This graph shows the number of photons per flash as a function of water temperature. At room temperature there are about a half-million photons per burst. Near freezing the amount of light goes up by a factor of 10. Curious.

[Transparency 25]

This graph (also from the ULCA group) shows what happens to the bubble as the water is cooled down. You would expect the dynamics to be a little bit different. However, the upper right graph showed that the equilibrium radius of the bubble does not change much as the water is cooled from room temperature down to freezing. The change in equilibrium radius of a part in 5, say, 20 percent.

The lower left graph shows the ratio of the maximum to minimum radius of changes by about a factor of 2. The acoustic pressure amplitude applied (upper left) changes by a few percent. All these parameters change, percents, factors of 2, but the amount of light (lower right) goes up by orders of magnitude. The light output is very sensitive to relatively small changes in parameters.

MR. WAKELAND: If you put some air in there and you turn this thing up and it goes through all these weird oscillations and finally there is a stable bubble in the middle, and then you do it again. Is the bubble the same size? Does it have the same average radius from experiment to experiment?

DR. ATCHLEY: No one knows how to control the ambient size of the bubble that I know of.

MR. WAKELAND: So, what was that that you just put up about radius?

DR. ATCHLEY: As measured.

MR. WAKELAND: So this is a single bubble?

DR. ATCHLEY: No, they are probably different bubbles.

MR. STOREY: For it to be in stable equilibrium, there is a size, so that the out-gassing and the in-gassing is the same. Otherwise, the bubble would grow or shrink. For a given set of parameters, there is likely a radius that is the stable radius. How you know what that is, is an unanswered question.

DR. ATCHLEY: Yes, you cannot control it, but you might believe that if you set the thing up the same time after time, the bubble is the same time after time.

DR. LAUTERBORN: It is different for different temperatures, for instance.

DR. ATCHLEY: Right.

About this time the newspapers got involved -- I do not think this is in your notes. I will just show some titles of articles: "Tapping the Light Fantastic," "Inferno in a Bubble," and "Sonoluminescence and the Heimlich Effect."

The 21st century talked about pursuing the cold fusion genie and they tied sonoluminescence to cold fusion.

Then someone uttered the F word, fusion, and started wondering if you can get neutrons from these bubbles. Are there thermonuclear reactions? This article is from The New York Times, and this one is from Science News. The press really got hold of this, because of all these hard to believe claims about this little bitty system with all these amazing parameters.

It really got bad when Hollywood got into it. They made a wonderful cinematic contribution called "Chain Reaction," in which Keanu Reeve plays Tom Matula -- (Laughter) -- and shows what can go wrong when graduate students are left in the lab by themselves.

It was a movie where several square miles of downtown Chicago were obliterated because people were -- the movie used the term -- "tapping the power of hydrogen" from oscillating bubbles. They tapped too much and blew up downtown Chicago.

Of course, the federal government was involved, because they did not want this -- I should stop talking, because it implies I saw the movie and I remember something about it. But rent it. If you have nothing else to do, rent the movie.

[Transparency 26]

The early research involved water and natural air. Natural air can be nasty, especially if you work in Los Angeles. The air that the UCLA people had to deal with was just terrible.

What if you make your own air? What is air? It is mainly nitrogen. Therefore, you put nitrogen in the water and you find that a bubble hardly glows at all. That is shown on the lower line of the graph. This graph shows the sonoluminescence intensity, normalized to the intensity of luminescence from an air bubble as a function of the dissolved gas pressure. The graph indicates that if an air bubble is 1, a pure nitrogen bubble in water is 20 to 50 times dimmer.

What else is in air. Oxygen. Air is about 80% nitrogen and 20% oxygen. So you mix some oxygen in with the nitrogen and the result is about the same, not much light.

What else is in air? Thanks to Lord Rayleigh, we know that air is about 1% argon. So, you mix 1% argon with 99% nitrogen. The intensity is almost the same as air. Now, if 1% argon is good, then more argon must be better and so you mix 10% argon. The intensity gets dimmer. You try pure argon and it is dimmer still.

It turns out that it was fortunate that Felipe was living in clean Mississippi, where the argon concentration was not obliterated by carbon monoxide and other stuff that you might find in Los Angeles.

DR. GARRETT: And it was already fully saturated with water. (Laughter)

DR. ATCHLEY: This graph also shows that the intensity depends on the partial pressure with which the liquid is saturated with that gas. I am not going to be able to get into this very much. The point is that the amount of light that is given off is very sensitive to not only the type of gas that is involved but also the concentrations of the gas.

[Transparency 27]

This graph from the UCLA group shows the optical spectrum of the emission. The previous graph was the intensity, how much light is coming out. This shows the spectrum for different types of gases.

The lines correspond to pure argon, 1% argon in nitrogen, 1% helium in nitrogen, 2% xenon in nitrogen, pure xenon, and pure helium. What is interesting about this is for 1% argon and nitrogen there is no real indication that the spectrum is going to peak anywhere. With pure xenon in the bubble, you start to see a peak in the spectrum. Maybe now a more complete understanding is starting to emerge.

[Transparency 28]

The most sonoluminescence-friendly liquid that we know of is water. It is difficult to generate sonoluminescence in liquids other than water, a fact that is also a little bit confounding. Multiple-bubble sonoluminescence has been known for a long time. Sonochemists love to use things like peanut oil and silicon oils and their systems glow very brightly. If you try to do single bubble sonoluminescence in those liquids, the results are dismally disappointing. Either you do not see anything or it is extremely difficult to see anything.

Water seems to be the friendliest liquid, although you can get sonoluminescence in other liquids. The point of this graph is to show that it is possible to generate single bubble sonoluminescence in other liquids, but the intensity tends to be lower than for water.

[Transparency 29]

The UCLA group initially reported an isotope effect when they compared sonoluminescence of hydrogen and deuterium bubbles in water and heavy water. However, they also found out that the spectrum was as a function of time. Over a period of days, the spectrum shifts significantly. This leads to the suggestion that some impurity is being introduced. You have to be extremely careful with your system to make sure that you know what is in that before you report any results. Purity is extremely important when you are going to start worrying about gas concentrations to such small percentages.

[Transparency 30]

A question came up earlier about emission lines I said then that one characteristic of emission from multiple-bubble sonoluminescence is the presence of spectral lines that depend on the nature of the host fluid.

[Transparency 31]

Tom Matula and his collaborators compared the spectra of single-bubble and multiple-bubble sonoluminescence. This graph shows results for salt water. The spectrum of multiple-bubble sonoluminescence shows emission lines from sodium and hydroxyl radicals.

For single-bubble sonoluminescence in the same fluid, there is no evidence of those emission lines at all. This is also something to be explained: Why is the spectrum so different in the two types of sonoluminescence?

DR. FORTUNKO: Is the physics of multiple bubbles understood?

DR. ATCHLEY: No. The parameter range is also different. Correct me, Tom; the multiple bubble sonoluminescence is done in a gassy liquid, right? Saturated with air?

DR. MATULA: Yes.

DR. ATCHLEY: Yes, this was air-saturated liquid. That is characteristic of multiple-bubble sonoluminescence; you do it in gassy liquids. In single-bubble sonoluminescence, you degas, so the system is a little bit different. The dynamics of MBLS are not understood any better than they are for single-bubble sonoluminescence.

MR. WENDOLOSKI: Are there any experiments somewhere in-between, like several-bubble sonoluminescence, 2 bubbles, perhaps?

DR. ATCHLEY: You can levitate bubbles at more than one point in a standing wave. People have done measurements. I do not think anything startling has happened.

DR. LAUTERBORN: We have done 2 and 3 and they blink alternately exactly as the sound field is, one and then the other.

MR. WENDOLOSKI: So they look just like copies of one, not like this other phenomenon at all.

DR. LAUTERBORN: Yes, because they are pretty far away.

DR. ATCHLEY: In MBSL, the bubbles are close together. They interact and either coalesce or eject one another.

Now I want to spend a few minutes talking about some more parameters. This is going to be necessarily incomplete, but I will point out some interesting aspects. There has always been a problem with stability of bubbles.

[Transparency 32]

Look at Figure 1 and note the pressure scale. Previously we talked about saturation pressures of hundreds of millimeters of mercury. Now we are talking about saturations of less than 10 mm Hg. A pure argon bubble behaves better at low pressures than at high pressures.

I had painted this picture of a very stable, very synchronous phenomenon. However, it turns out if you change parameters a little bit, then you can see all sorts of cyclic phenomenon.

[Transparency 33]

This graph is from a paper by Bob Hiller. It shows the intensity of light as a function of time. The time scale is seconds. The sonoluminescence intensity is low for a second or so, which is a lot of acoustic cycles, then it jumps up and is high for a while and then jumps back down. This is a long-term phenomenon that happens over lots and lots of acoustic cycles.

[Transparency 34]

Felipe and Glynn Holt, Joachim Holzfoos and I looked at the effect of slightly de-tuning the resonator. If the resonator is driven on resonance, the bubble is very stable. If you de-tune the resonator just slightly, maybe 10 Hz out of a few 10s of kHz, you find that, instead of being extremely synchronous, where the arrival time between successive flashes is the same almost every time, you get 2 different arrival times. The flashes come a little bit earlier, a little bit later, a little bit earlier, a little bit later. With more de-tuning, you get 4 different arrival times, and then you get a whole smear of arrival times.

If you analyze the arrival times using dynamical systems analysis, you can see that you can have a period one motion (Figure a) that is extremely synchronous. You can have period-2 motion (Figure b), 2 arrival times; period-four motion (Figure c), and then you get to this quasi-periodic regime (Figure d).

[Transparency 35]

If you de-tune even further, the arrival times become chaotic. This picture of synchronicity is not always true. You can have a lot of different behavior, depending on the parameter range that you select.

I have spent a lot of time telling you about the behavior of the light, but I have had almost nothing to say about the behavior of the bubble. What is the bubble doing? Is the sonoluminescence behavior related to the bubble dynamics? If you see a change in the intensity of light, does that mean the bubble is behaving differently? Just what goes on?

[Transparency 36]

If you scatter laser light from a bubble, you find that more or less the amount of light scattered depends on the cross-sectional area of the bubble. There is a proportionality factor that is important, but bigger bubbles scatter more light, smaller bubbles scatter less light. If you can measure the intensity of the scattered light, you can infer something about the size of the bubble.

This graph is more of Brad Barber's work. The dots are averages over many cycles of the light scattering from the bubble as a function of time and the solid lines are fits to the Keller equation (or maybe they used the Gilmore equation). In either case, it is a fit to a hydrodynamic theory.

I will not go into the details of what parameters they considered free parameters. There is a lot of discussion about their letting viscosity and surface tension be free parameters. The point is that you can fit the radius-versus-time curves, to data well. These fits accurately track the bubble behavior over almost all of its cycle.

The trouble is that sonoluminescence occurs right in a very small fraction of the cycle, at a time when the bubble is undergoing extremely rapid changes. The question is, what happens in those areas that are tough to probe optically?

[Transparency 37]

I should say there are other techniques to measure radius-versus-time curves. Tom Matula and his group have developed methods where you do not have to average over cycles, you can do instantaneous R.T. curves. Professor Lauterborn and Claus-Dieter Ohl and Bob Apfel's group use photographic methods.

Examine the upper curve. The bubble hangs around near its equilibrium size; it grows, collapses quickly, has big bounces, and then repeats the process. It turns out that at this drive amplitude this bubble is not emitting light. If you increase the pressure amplitude a little bit more, you see a sudden change in the behavior of the bubble (lower graph). The maximum radius gets smaller and the big bounces disappear. After this transition occurs, light begins to appear. This graph shows the transition as well. The heavy line is the intensity of the emitted light. So there is this transition in the behavior of the bubble and then light turns on. Then, as you drive the bubble more, the maximum radius gets bigger. You get more and more light. Finally, you get to a point where if you drive the bubble harder and it disappears.

This transition looks like a pretty neat clue. Now you can tie bubble dynamics to some light-emission mechanism. The question is why do these bounces disappear when the light starts to come out?

[Transparency 38]

This sort of threshold phenomenon looked pretty interesting. Then they did it with a pure argon bubble in water, and they did not see this threshold phenomenon. This bubble glowed the whole time and the bounces were there the whole time. People are still trying to understand why there is a difference in air bubbles and argon bubbles.

[Transparency 39]

Researchers at Yale have put surfactants in the water to see if they can affect the dynamics of the bubble. If a surfactant is attracted to the bubble surface, perhaps it changes the bubble dynamics, and perhaps the light intensity.

They tried 2 different surfactants, one, triton X, which is a surfactant that is known to not impede interface motion very much, and another one, BSA, which is bovine serum albumin, which is known to affect, impede, interfacial motion. They find that the addition of surfactants changes the bubble dynamics, the characteristics of the light emission, and how hard you can drive the bubble.

[Transparency 40]

Researchers have also studied sonoluminescence in high magnetic fields with interesting results. Figure 1 shows the amount of light given off as a function of the applied magnetic field for different acoustic drive amplitudes. They show that if you increase the magnetic field, you

quench the luminescence. You can make it turn off. The measured curves can be fit to the square of the magnetic field.

They also found (Figure 3) that the presence of magnetic field alters the stability range; that is, the range of pressure amplitudes over which you can get sonoluminescence. There is also a dependence on water temperature.

DR. GARRETT: Do you have any feel for the Maxwell stresses at those levels?

DR. ATCHLEY: I do not have a feel for it.

DR. GARRETT: Are they just modifying the pressure of the liquid or are they --

DR. ATCHLEY: They throw out a couple of conjectures. One is that you are changing the stress in liquid or, actually, if there is a plasma in the bubble, then maybe you are concentrating magnetic field lines. So as the bubble collapses a stress is induced that balances the acoustic pressure

[Transparency 41]

One more topic before we take a break. Of the 50 or so papers in the last couple of years, I think this one has to be one of the most important. This paper is by Glynn Holt and Felipe Gaitan and I really urge you to read it. It is well written and very interesting.

They mapped out a little bit of parameter space. They took bubbles and tried to map out a space of equilibrium radius as a function of pressure amplitude for different gas concentrations. They find that if, for instance, you have a bubble of this size equilibrium radius, which is, I guess, about 15 μm , it is stable at a certain pressure amplitude. That is, you can levitate it stably at a pressure amplitude of about 0.8 bar.

The solid line is a line of stability predicted from a simple theory of gas diffusion by Eller and Flynn. If you drive a 5- μm bubble at 3/4 bar pressure amplitude, that bubble will dissolve. It will not be stable; you cannot counteract the diffusion with the acoustic field.

If you drive the bubble at around 1 bar, then you can make it stable against diffusion. If you drive it harder, that bubble will grow. At some point, if you drive it hard enough, you get sonoluminescence. What is interesting is that the sonoluminescence in this parameter space occurs only along a line (the dashed line in the lower figure). The colors do not show up so well. Sonoluminescence occurs only along a line of pressure amplitude versus equilibrium radius, a line in this parameter space.

The extent of that parameter range depends on the gas concentration. Along this line you find bubbles that are stably levitated and emitting light. There is a process called rectified diffusion that is important. As the bubble expands, the pressure inside the bubble is small. Therefore, if the bubble started off in diffusive equilibrium with the surrounding liquid, then the concentration of gas in the bubble is smaller when it expands, and gas diffuses into it.

When the bubble shrinks, the concentration is higher and the gas diffuses out. It turns out that the diffusion of gas in this range is rectified so that more gas flows in than flows out. If you drive this bubble cycle after cycle after cycle, it pumps itself up and it keeps growing. This process is called rectified diffusion.

The growth rate and equilibrium size depend on the concentration of gas in the liquid. It turns out that this (solid) stability line obeys simple gas diffusion, relatively well. Along the dashed line, the stability line for sonoluminescence, they had to use a gas concentration that was a factor of 100 lower than what you would guess from the concentration of air you put in the

liquid. That is, you have to use a much lower concentration than what you would nominally call the saturation concentration of the air in order to fit that line.

The question is how come? What has happened to all of that gas and how come this stability appears to behave as if the gas concentration is much lower? If I had to put money on it and I had to tell you the most popular theory of sonoluminescence, I guess that is a big clue. To find out where that clue leads, you have to come back after about a 5-minute break.

DR. ATCHLEY: What I have tried to do so far is to point out some of the things that are known experimentally and also to ask if there is any way of understanding what is going on.

[Transparency 42]

I put this picture up just to remind us that we have a system that is undergoing extremely rapid and violent collapse. You would like a consistent theory of how the bubble wall is driving what goes on inside the bubble.

People find that under a certain range of conditions, that when the bubble collapses a shock wave is launched into the gas ahead of the bubble wall. This shock wave compresses the gas inside the core of the bubble.

Since the test ban treaty was put into play, the employees of Los Alamos and Livermore had nothing to do. So they said, well, what else do we know that gets hot and gets compressed. Sonoluminescence! They have some pretty sophisticated codes and they know something about hot gases and extreme conditions. If you couple bubble dynamics with real equations of state, a wealth of information comes out about shock-wave dynamics inside the bubble.

The question is, can a shock-wave model explain sonoluminescence. Can it explain the short pulse time? It can do better than a non-shock-wave model. An adiabatic heating model should compress and expand the bubble on hydrodynamic time scales. The pulse lengths are much shorter than hydrodynamic time scales. Shock wave can help with that, because things happen only when the shock wave converges on the center of the bubble. It collapses and reflects over very short time scales.

Can shock waves models explain some of the other dependencies, for instance, the temperature dependence of the water? Why should changing the temperature on an absolute scale a pretty small amount change the properties of the light so dramatically?

[Transparency 43]

Some people have looked at the water-temperature dependence of single-bubble sonoluminescence. This paper came out in Physical Review Letters in February. They say that if you are careful about keeping track of the properties of the water, -- the surface tension, the viscosity, and the vapor pressure inside the bubble -- you can explain a lot of the dependence on water temperature from shock-wave models of sonoluminescence.

[Transparency 44]

Are there shock waves? Here are some results from Tom Matula. Figure 3 shows the radius of the bubble versus time determined from laser scattering and the output of a hydrophone positioned some distance away from the bubble as functions of time. When the bubble collapses, shock waves are radiated into the fluid. Therefore, shock waves do exist, outside the bubble anyway.

DR. FORTUNKO: What is the fingerprint that you are using to tell that it is a shock wave?

DR. ATCHLEY: How do you know it is a shock wave? I will pass.

DR. MATULA: We do not know. That is actually from a focused transducer sitting a long way away from the bubble. If you take a hydrophone that is calibrated and you stick it very near the bubble, say a millimeter away, you measure a pressure that is on the order of a couple of bars.

If you assume the bubble emitted this wave when it was about a micron, the hydrophone is 1000 times away; so just geometrical spreading itself would make the pressure at the bubble about 2000.

DR. ATCHLEY: There is also photographic evidence that we will get to.

There is evidence that shock waves exist. What about the pulse duration? What is the pulse duration? The best estimate was based on resolving rise times. I said earlier that one of the most important papers in the last 2 years is by Holt and Gaitan that explored parameter space.

[Transparency 45]

I think the other one has to be this paper by Gompf and colleagues in Germany. They used time-correlated photon counting to resolve the pulse duration of sonoluminescence. They found that the pulse duration of the emission is on the order of a few hundred picoseconds, so maybe 200 psec for certain parameter ranges.

They also investigated how the pulse duration depends on drive amplitude. Although there is some variation, the pulses are somewhere in the hundred to a few hundred picoseconds range, full-width half-max.

They also looked at the pulse duration for different colors of light, different parts of the spectrum, and found that it is the same. The pulse duration is largely independent of the spectrum. The pulse duration in the red is the same as the pulse duration in the blue. This result can largely rule out adiabatic heating, because if it is a slow compression you would expect the emission would turn on in the red before it would turn on in the blue. This paper is important, not only because it resolves the pulse duration, but also shows that the pulse duration is independent of the spectrum.

What does that tell you? It leads some people to believe that the light emission mechanism is caused by the shock wave forming a plasma in the gas. This plasma quickly turns on and turns off. It is free electrons that radiate from bremsstrahlung.

[Transparencies 46 and 47]

I included these two papers to complete the references on pulse durations and pulse shapes.

Shock waves can tell you that you can get fast pulses. In addition, if a plasma forms the pulse duration is independent of spectrum. It can also explain the temperature dependence, maybe, by taking into account the real properties of the fluid as a function of temperature.

The really perplexing thing left to explain is the gas concentration dependence.

[Transparency 48]

An interesting hypothesis came out in February of last year. It says that if the bubble starts off filled with air and is driven at low amplitudes, it undergoes nice linear volume pulsations. Nothing very dramatic happens inside.

However, if the amplitude is increased, extreme pressures and temperatures are generated inside the bubble. This extreme environment causes the dissociation of the nitrogen and oxygen, and all the molecular constituents of the gas. They form free radicals, highly chemically reactive species that get forced out of the bubble and into the liquid.

It is only the inert constituents of the gas that do not dissociate and remain inside the bubble. Therefore, even if you start with a bubble of air, what is left after some number of cycles is argon. The nitrogen and oxygen have gone out and you have an argon bubble.

This explains why the data from Holt and Gaitan showed that for the region of stability for sonoluminescence, the gas concentration had to be so low. The value they had to use pretty much matches what the natural constituency of concentration of argon and water is in the first place.

[Transparency 49]

The next paper by Tom Matula addresses this issue. If argon rectification is true, then the following ought to be true. First off, pure argon bubbles should behave differently from air bubbles, because argon bubbles have argon, to start with.

Maybe more interesting, if I take an air bubble and I drive it into the parameter range where it emits radiation, sonoluminescence, it has to go through this argon rectification the first time I do it.

Now suppose I have an air bubble that has been radiating sonoluminescence for a long time. If I take it below the sonoluminescence threshold for a while (not too long, a few seconds) and then I turn it back up, the argon should still be in there. The other gases have not had enough time to diffuse back yet to any great extent. Therefore, the evolution of the light intensity for a first-timer should be different that for one that has been there before.

They find exactly that, that there is a difference in the behavior. This that would tend to indicate that there is evidence for gas exchange in the bubbles.

DR. GARRETT: Has anybody looked at the rebound? Before you go to sonoluminescence, you get this nice rebound. Then the transport properties of the gas mixture that have not been purged are different from the transport properties of the other, so the damping of the bubble should also be dependent on whether this chemical thing happens. Maybe those nice Hiller data where you see the rebounds go away would be confirmed.

DR. ATCHLEY: Tom Matula says the reasons the rebounds go away when you get to the sonoluminescence state is that the shock wave takes away a lot of energy from the bubble. Therefore, there is not enough energy left for the rebounds.

DR. GABRIELSON: I was wondering if anybody had looked in the early stages of that process for a nitrogen peak in the spectrum.

DR. ATCHLEY: That is pretty interesting. The light is pretty dim then. That is a good question. I do not think so. You would have to catch the spectrum on the fly.

I could almost stop, but I will not, because I have 20 more minutes. I have presented a picture that says everything is well understood, that all the evidence is that you have spherically symmetric bubble oscillations, spherically symmetric shock waves that lead to a plasma generation in the center of the bubble and that all these weird dependencies can be taken into account by a good calculation of the bubble dynamics and the chemistry. Argon rectification ties it all up in a nice package.

[Transparency 50]

"Sonoluminescence from an Isolated Bubble on a Solid Surface." That bubble is hemispherical and it emits light.

[Transparency 51]

Here are some photographs from Professor Lauterborn that show high-speed photography of collapsing bubbles. Take a look at this. It shows a bubble collapsing near a solid surface. It undergoes a radial expansion, then it collapses and rebounds. The collapse is clearly an asymmetric.

[Transparency 52]

Here is a blowup of one of those frames -- it has got to be one of the neatest pictures in fluid dynamics. There is a lot of evidence that bubbles collapse asymmetrically.

[Transparency 53]

Here is a picture from Larry Crum that shows a bubble collapsing on a solid surface. All these pictures have surfaces involved. The reason that you get asymmetric collapse in the presence of the surface is because the fluid flow around the bubble is asymmetric. Therefore, you get this distortion and a jet that penetrates through the bubble. If the bubble is on the surface, the jet hits the surface and causes erosion. If the bubble is far from the surface, the jet goes through the bubble and hits the other side of the bubble.

Now, there is not a solid surface near the bubbles that I have been talking about for the past 2 and a half hours. So, is there any reason to believe that there should be an asymmetry? The answer is yes. If you take a bubble that is levitated far from a boundary but in a standing wave, it is not levitated exactly at the pressure anti-node; it is displaced because of buoyancy. Therefore, it experiences a pressure gradient, so the sound field around it is asymmetric.

Also, I sort of lied to you -- not sort of, flat out lied to you -- when I told you about levitation, because we did a time average and we think in terms of a static force. Actually, there is also a dynamic force on the bubble that varies with the acoustic frequency. This force makes the bubble translate up and down. The bubble experiences a pressure gradient and oscillates up and down. There is a lot of reason to believe that the bubble should collapse asymmetrically. In fact, some people believe that you cannot have a symmetric bubble collapse.

Is there a way that you could have sonoluminescence from an asymmetric bubble collapse?

[Transparency 54]

These pictures are downloaded from Claus-Dieter Ohl's home page. You should download this and look at it in the video mode. It shows a bubble collapsing asymmetrically. It becomes very distorted. What you cannot see so well are the details of the shock-wave formation, but there are shock waves given off by this asymmetric collapse.

[Transparency 55]

Here is another result from the Lauterborn group. It shows a bubble that collapses asymmetrically. It is toroidal in shape. A jet forms and the bubble is highly distorted. Three different shock waves are coming off this bubble for three different reasons. One is because the jet is forming, another is because the bubble is collapsing around the torus, and the third shock wave is given off when the bubble reaches its minimum size. There are lots of ways to get shock waves and complicated dynamics from asymmetric bubble collapses.

[Transparency 56]

There are hypotheses for sonoluminescence from asymmetric bubble collapses. I will talk about a couple of them just for completeness. Longett Higgins does a lot of calculations of surface instability of bubbles. When a bubble collapses, not only do jets form, but surface distortions and instabilities also occur.

Some people say that these surface instabilities lead to a pinching off of the bubble surface, producing droplets inside the bubble. When these droplets pinch off electrical charge is separated, leading to high electric fields that produce sparks.

[Transparency 57]

Andrea Prosperetti claims that you cannot have a symmetric collapse. When the bubble implodes, a jet starts to form on one side of the bubble. It penetrates the bubble and hits the opposite wall when the bubble collapses. If the jet is fast enough, it hits the bubble wall on time scales that are too fast for the water to behave like a liquid. It reacts as a solid. The jet is so energetic that it cracks the water and it is the fracture of the water that leads to the sonoluminescence.

[Transparency 58]

The next article is by Claus-Dieter Ohl, Professor Lauterborn and others. These photographs did not show up at all well on the copies. This paper tells us that you can get luminescence from both symmetric collapses and asymmetric collapses.

They used a focused, pulsed laser to form bubbles in the liquid. They can position the laser any distance away from a solid boundary. Therefore, they could control the symmetry of the collapse by where the bubble was formed. If it is far away, the bubble collapses nearly symmetrically and light comes from the center of the bubble.

As the bubble is formed closer to the wall, the collapse is more asymmetric. These bubbles also emit light.

As the bubble gets closer to the wall, the asymmetry gets even stronger but then the emission disappears. Therefore, some asymmetry is good but a lot of asymmetry is bad. What does this mean for Prosperetti's hypothesis?

[Transparency 59]

Here is an interesting paper published in Physical Review Letters a couple of years ago. It discussed the transition from normal to fast sound in liquid water. The horizontal axis is the momentum of the particle of the water molecules in units of inverse meter. The important point is that as the momentum of the water molecules increases, you observe a change in the sound speed.

At low momentum, it is 1500 m/sec that is what we think the sound speed in water is. At high momentum, the sound speed is higher, something like 3200 m/sec.

Prosperetti might argue that if the bubble is close to the boundary, then even though the bubble is highly distorted on collapse, the jet motion might be too slow to fracture the water. There ought to be a parameter range, depending on the jet velocity that the jet hits with enough momentum that the water reacts like a solid and cracks. You might surmise that the dependence of the luminescence on how far the bubble is from the surface is explained by jet velocity.

DR. MIGLIORI: What is the speed of sound in water?

DR. ATCHLEY: Fifteen hundred meters per second.

DR. MIGLIORI: So they got one of them right.

[Transparency 60]

DR. ATCHLEY: What else could happen? People are dying to get quantum mechanics involved with sonoluminescence. This is a paper by Claudia Eberlein that says that sonoluminescence is quantum vacuum radiation. I am not going to go into any of the details.

This paper received much criticism. Regardless of whether or not it is right or wrong, this paper took a beating. It appeared in May of 1996 and over the next few months, people piling on.

[Transparency 61]

After all the criticism, this article appeared in March of 1998. The author writes, "The similarity of the effects of highly resonant, non-thermal rf. . . and sonoluminescence suggests that quantum vacuum radiation might indeed be the most feasible model to explain both phenomena."

You can take whatever answer you want away from this talk. You could have stopped listening to me 20 minutes ago and believe it is symmetric shock waves. You could have listen for another 10 or 15 minutes and believe it is really asymmetric collapses and cracking water. Or you can listen to the last thing I said and not know anything about sonoluminescence.

With that, I think I will conclude and thank you for your attention and I invite people who really do this for a living to stand up and say what they think is going on. Thank you.

DR. HOFER: In the experiments with the asymmetry during the collapse, you can control it so it becomes more asymmetric earlier in collapse, I guess. Then you correlated that to the jet velocity. I am wondering --

DR. ATCHLEY: I hypothesized.

DR. HOFER: Hypothesized, right. I wonder to what extent, if you had more asymmetry earlier, that you would end up with perhaps less pressure in the bubble, so maybe the overall pressure inside the bubble when the jet comes through is -- would it matter?

MR. OHL: You should look at the last stage of bubble collapse and you should really define what is symmetric collapse and what is asymmetric collapse. Incidentally, we have never seen a symmetric collapse. The asymmetry starts very, very late. This means you are far away from the boundary and the asymmetry starts very late. You have very high speed, where the bubble walls impact each other. The boundary is very far off. The influence is there but it is small. Then you have a high-speed impact.

DR. HOFER: Do you measure or anticipate any general difference in volume during the jets in the one case versus the other? I mean, is it a smaller bubble when it is later jetting as opposed to the earlier?

DR. OHL: The size of the light-emitting region is about the same.

DR. HOFER: So you do not have any reason to think the pressure would be different?

DR. LAUTERBORN: -- minimum bubble radius or minimal form. Nobody knows what that is. That is a problem, one of the problems.

DR. ATCHLEY: It is really hard to know in that last instant, where all the neat stuff happens, what is going on. You cannot probe it optically very well, you cannot really put a probe in there to see what is inside the bubble. Do you want to show your tape in the last few minutes, Tom? [Video demonstration]

DR. MATULA: We used a red laser and strobed it against the bubble. You can get it to grow and collapse. This is sonoluminescence, but the little light in the center is probably just the light coming through the bubble. The laser is on the other side and you are looking at the shadow.

If you look closely, you might see that the rebounds are a little bigger.

PARTICIPANT: What is the fluid?

DR. MATULA: Air and water.

The drive frequency is about 25 kHz.

DR. ATCHLEY: The next sequence is the shock wave. If you look really quickly, you can see a shock wave radiating out from this bubble.

MR. STOREY: I have a comment about the whole shock wave model. It would be especially important if it is the argon rectification, because with argon you have a high ratio of specific heat, so you have a lot of adiabatic compression long before you would launch a shock wave.

The problem with developing a shock wave in that case is the center is very, very hot. It is going to be difficult for the pressure wave in the gas to converge into a shock wave, because you have got this larger amount of adiabatic heating long before the bubble starts to collapse. You are always going into this adverse region, so it will not steepen into a shock.

But if you do the same thing and you assume all the same parameters and admit some mysterious gas that now has a specific heat like air, you do not get as much heat and so you do get a shock wave.

DR. LAUTERBORN: I think Moss also has calculated that.

MR. STOREY: Yes, you have to include the diffusive terms. You also have to include normal heat conduction and normal viscosity. A lot of the calculations that Moss did pursued these plasma equations. If you just assume from the beginning that you have heat conduction and viscosity and you never have a shock wave, you never get the intense temperature you need for the exotic plasma equations.

SENSOR PHYSICS: SIGNALS AND NOISE

Thomas B. Gabrielson
Applied Research Laboratory
Pennsylvania State University

DR. GABRIELSON: A few months ago, Gunnar Rasmusson came to Penn State to talk about his work on precision condenser microphones. He opened his talk by saying that he had just finished reading an article about the proliferation of data in recent years. The title of the article was something like "Separating Noise from Knowledge." Gunnar said that this was nonsense: noise is knowledge. That summarizes the first two hours of my talk. (Laughter)

Noise is definitely knowledge. The more you know about noise, the more you know about your system. It is appalling to see how often people design advanced sensors without understanding noise. As a result, they propose things that do not make sense. Even more unfortunate, some of the systems could have value if they were redesigned with some intelligence toward noise.

I will also talk about signals a little bit later, because the idea of studying sensors has two very interesting aspects. One is the aspect of sensor noise itself. If you think sensor noise is connected with only the pre-amp attached to that sensor, you are wrong. There are some very interesting physical noise phenomenon associated with the sensor itself that can be important in some instances.

Another area that is fascinating is finding out how well sensors respond to signals, not just in a crude sense, not just an accelerometer on a shaker and comparing it to a reference, but trying to get the best possible measurement of response, finding out what all those little wiggles mean in the sensor response.

If you take these two extreme positions, you will learn wonderful things about that sensor and about sensor physics in general.

Part I: Equilibrium Thermal Noise

[Transparencies I-1, 2]

At first, I am going to talk about very simple concepts. I am not going to require hard math -- some of the concepts are a little obtuse and I will try to clarify those later -- but I am going to cover the basics first: equilibrium-thermal noise in mechanical, electrical, or optical systems. These are the simplest kinds of noise phenomena that we find. There are global principles that govern thermal-equilibrium noise and it is often quite simple to calculate.

Next, I will talk a little bit about shot noise and nonequilibrium noise and then I will talk about some techniques for precision sensor calibrations. That is where I would like to reintroduce the response to signal.

[Transparency I-4; skip 3]

The first question I would like to pose concerns the simple harmonic oscillator. Consider the ball on a spring as a single-degree-of-freedom system. There are very few true single-degree-of-freedom systems, but it is an extremely useful approximation. You can measure some

of the properties of a system like this. You can measure its resonance frequency; you can measure its Q , for example. (I've included some slides about Q [Transparencies I-5 and I-6] in your handouts.) Suppose we had a system with a resonance frequency of 100 Hz, a Q of 10, and a mass of a milligram. This could represent some prototype accelerometer that we wanted to build.

If you displaced the mass, if you gave this system an initial displacement of a millimeter, how long would it take that oscillation to decay down to 10^{-8} mm? You might argue, first, that it is ridiculous to talk about 10^{-8} mm, but be careful with your intuition. At 100 Hz, that corresponds to about a micro-G, and in some applications that is a rather large signal. We are not talking about a pathological case; we are talking about something relevant to real sensors.

Anyway, how long does it take? It is a very simple calculation. You can calculate the rate of decay (it's exponential) from the Q and the resonance frequency; from the rate of decay, you can figure out how long it takes. The only difficulty is that you would be wrong: the system never reaches an amplitude of less than 10^{-8} mm, on average, and that is because there is something forcing the mass to move even when you are not.

[Transparency I-7]

The spring-mass oscillator is made up of molecules, and these molecules are in constant motion. (I am talking about a system at room temperature.) There are also air molecules hitting the mass from the outside, but even if the system were in a vacuum, there would still be molecules in the structure and they would be vibrating and shaking the mass. In this particular system, the vibrations are strong enough that the oscillation never decays to the 10^{-8} mm level. Therefore, you cannot make a micro-G accelerometer out of a system like this without understanding this phenomenon and dealing with it.

DR. MIGLIORI: What is the lowest amplitude?

DR. GABRIELSON: I do not remember, but I will show you how to calculate it, which is even better.

[Transparency I-8]

There is a theorem that is well known thermodynamics called the fluctuation-dissipation theorem. It is a very nice theorem, because it has a very fancy name but you do not need any math to extract the real meaning of it. In fact, I have seen very few people actually use the mathematics underlying this theorem.

Let's look at a structure similar to that of a condenser microphone. There is a diaphragm suspended by some kind of flexure so that the diaphragm can move back and forth under excitation (an acceleration or an acoustic wave impinging on it, for example). The diaphragm is probably close to another surface and there is a gas between the diaphragm and the back plate.

What are the mechanisms of energy loss in this system? There are many mechanisms of energy loss here. One of the biggest ones in condenser microphones and in microfabricated structures that are operate at atmospheric pressure is the damping produced by the viscosity of the fluid between the moving diaphragm and the nonmoving plate. As they come together, these two surfaces squish the fluid out and as they do that, a boundary layer is formed and drag is produced. For very small spacings, which you have in condenser microphones or in microfabricated devices, these drag forces can be very high.

These systems can have Q 's of much less than one. The viscous damping can produce a very large effect. Even if you evacuate the system, there would still be damping associated with the mechanics of this plate itself, the flexures and the boundary conditions. There are always mechanisms of internal loss.

If the sensor is exposed to the environment, it can radiate: if you put the diaphragm in motion somehow, the diaphragm moves and it radiates sound. You do not usually call this radiation dissipation, but from the point of view of the system, it is an energy leakage away from the system, so from that point of view it is a dissipation.

Thermal radiation also takes energy out of the system. This is normally not significant unless you were building a sensor to operate at very high temperatures, but if you ever need to design a sensor for such an environment, you may need to consider this loss mechanism. At high temperature, the thermal radiation can take a significant amount of energy out of the system.

The fluctuation-dissipation theorem is really a statement that pathways for energy transmission work both ways. If there is a path by which energy can leave the system, then there is also a path by which energy from the environment can come back into the system.

Sometimes the energy that leaves the system is ordered energy ("signal") and sometimes it is disordered energy ("random vibrations" or "heat"). In the case of thermal conduction, the energy is disordered energy. If it is acoustic radiation, it may be ordered energy, but the energy that comes back from the environment is typically disordered energy. The energy reaching the sensor along the return paths acts as a forcing function and it forces the mass or diaphragm to move. Whenever a system has dissipation, it also has fluctuation associated with this dissipation. If the system is in thermal equilibrium, calculating the degree of fluctuation is simple.

[Transparency I-9]

One of the nice things about equilibrium thermodynamics and fluctuations is that you can find out how large the fluctuations will be rather easily. The equipartition theorem says that each degree of freedom of a system settles to this energy level: one-half Boltzmann's constant, K_b , times the absolute temperature, T . The value $K_b T$ at room temperature is about 4 times 10^{-21} joules.

If you have a point molecule in free space, it has 3 degrees of freedom. It can go up and down, left and right, forward and backward: there are the 3 independent coordinates.

For each of those coordinates, there is a kinetic energy associated with the motion: $1/2$ mass times x -velocity-squared, $1/2$ mass times y -velocity-squared, and $1/2$ mass times z -velocity-squared. Each of these energies settles to the $1/2 K_b T$ value just because the molecule is in contact with the environment at this temperature T . Since we know how much energy is associated with this motion, you can predict what the RMS velocity is in a particular direction.

One of the important things to understand is that it does not apply to just molecules; it applies to any object. It applies to a ball, if I throw it up in the air. The difference is that the mass would be a lot larger. The mass is a lot larger, so the RMS velocity is going to be a lot smaller, but it is not zero; it is predictable by this expression.

There are many of other kinds of systems. Consider an atom stuck inside a crystalline lattice. You can treat the atom as if it is connected to the lattice with little springs. There is a potential energy associated with compression or extension of these springs. If you can calculate that potential energy, you can express it in terms of the thermal energy. You have all of this

calculational apparatus to figure out what the energies of the system are and each type of energy settles to $1/2$ K_BT.

DR. FORTUNKO: I find that most people do not understand is that these energy quantities are really time averages. That is a difficult concept for most people to understand.

DR. GABRIELSON: Yes, your point is that the velocity, for example, is not exactly equal to the value given by these expressions but that the average velocity is equal to that value. The average can be a time average for a single particle or it can be an "ensemble" average (an average over many particles).

DR. MIGLIORI: I was just going to make contact with what Dr. Maynard said last night in that you have $1/2$ K_BT per accessible degree of freedom. For example, in a solid at low temperatures, even though there are $3n$ degrees of freedom, you do not have a specific heat consistent with that.

DR. GABRIELSON: Yes, they have to be degrees of freedom that you can actually excite with this level of energy. As you raise the temperature, you can see the $1/2$ K_BT quantity increases but you still need at least one quantum of energy to excite a particular state. If $1/2$ K_BT is less than the quantum value for a particular state, then it is not excited at all. A vibrational or rotational state for an atom in a solid may have a large enough lowest-level energy that $1/2$ K_BT cannot excite it at room temperature. Classical physics says that you would excite such a state a little bit, quantum mechanics says that you do not excite it at all.

DR. MIGLIORI: And that was a big puzzle in the specific heat of solids. Basically, that mystery was solved via quantum mechanics, why a solid has very low specific heat at low temperatures.

[Transparency I-10]

I've tabulated some energy levels here to give you some perspective. I am not going to use joules for units; I am going to use electron volts. I already told you what K_BT at room temperature was in joules. In electron volts, it is 0.025 electron volts.

Do not be concerned too much with the units. I want you to compare the numbers from one case to another. The quantum for visible light, hf , is in the neighborhood of $1-1/2$ to 3 eV, depending on where you decide the boundaries for visible light are -- my boundaries are gradually shrinking with age.

If you calculate the energy associated with an acoustic wave at 100 micropascals in air (a fairly low-level but certainly easily discernable acoustic wave), you get a volumetric energy density of about 0.4 eV/cubic centimeters. Reduced to the volume of a single atom, there is very little energy to excite that atom.

An acoustic wave in water with the same pressure level produces an energy density far lower: 20 micro-eV/cubic centimeters. Consequently, a very strong acoustic field will not excite interesting chemical reactions unless something other than uniform waves are involved. In sonoluminescence, something different is happening. There is a collapse in energy density from the volumetrically averaged acoustic wave down to a very small volume and the energy is concentrated enough to produce light.

For chemical bonds, covalent bonds can be broken at about 4 eV and ionic bonds can be broken with a bit less energy than that. Hydrogen bonds are responsible for many of the

interesting properties of liquid water and of biological systems. These bonds can be broken at about 0.2 eV. Why is this important?

This is important because it is a lot closer to $1/2 k_B T$ than a covalent or ionic bond, so that water has a lot of features that are much more temperature-dependent than other kinds of chemical systems.

[Transparency I-11]

Let's take a look at some specific sensing systems. The protozoan in this slide is a fascinating creature. The overall length is about 200 micrometers. Near one end (the top in the slide) is a small chamber, about 7 micrometers in diameter; inside that chamber (which is fluid-filled) is a hair structure with a lump of barium sulfate on one end. It looks just like an accelerometer. You can picture moving this animal up and down and this hair structure would lag behind because of the mass on the end of it. If the animal is pointed up, the barium sulfate drags the hair downward; if the animal is pointed down, the hair is dragged in the other direction with respect to the animal's body.

This creature likes to live at a particular level of oxygen concentration in a lake. Oxygen concentration varies from a maximum at or near the surface to a minimum toward the bottom, and it thrives at an intermediate level of oxygen concentration. It uses the accelerometer-like structure coupled with an oxygen sensor to figure out how to get there. If the ambient oxygen concentration is too low, the creature moves upward; if the concentration is too high, the creature moves downward. The actual strategy is more complex and fascinating in its own right, but I'm not going to discuss details of the strategy.

The sensitivity of this gravity sensor is adequate to detect a one-g acceleration but without much margin. If the linear dimensions of the sensor were reduced by a factor of 4 (i.e., the mass were reduced by a factor of 64), then the $1/2 k_B T$ value would equal the change in potential energy of the sensor mass from one extreme to the other. Even a reduction by a factor of 2 would probably not leave enough signal margin for reliable operation.

DR. MIGLIORI: If I can signal-average, can I extract the right answer because the gravitation potential, mg , is coherent in some sense?

DR. GABRIELSON: Normally, yes, but in this case the protozoan is constantly moving and changing its orientation with respect to the gravity vector somewhat randomly. This limits the effectiveness of averaging to time intervals short with respect to natural rotation of the creature's body.

DR. MIGLIORI: If $mg/k_B T$ is 330 for example, it tells you how quickly it can tell which way is up.

DR. GABRIELSON: Exactly.

[Transparency I-12]

If you thought that was interesting, here is a smaller version. This is a bacterium. This creature is about 3 micrometers long. It could fit inside the gravity sensor of the protozoan. It's far too small to use a mechanical gravity sensor to distinguish up from down so how does it work? Magnetically. It accumulates particles of magnetite that are small enough that they consist of single domains. That means that they are naturally magnetized.

The particles line themselves up and it turns the bacterium into a compass needle covered with protoplasm. You can argue that the protozoan is pretty clever in the strategy that it uses.

The bacterium is not terribly smart at all. All it wants to do is to get back down into the sediment after the sediment has been roiled up by a storm, so it just goes back down.

Now why would you use a compass to go down? This bacterium is found in northern latitudes. There the magnetic field lines dip down fairly steeply and the bacterium does not care if it is displaced horizontally. You do not want to think 2-dimensionally when you think about the earth's magnetic field. In the southern hemisphere you find this same organism with the granules of magnetite oriented the other way.

You can also, if you are not worried about the animal rights activists, put this in a large magnetic field and flip all the domains and put it back in the water. Then it goes up instead of down and starves. (Laughter)

[Transparency I-13]

This $1/2$ K_BT idea is a very useful idea, but what is even more useful is how this noise energy is distributed over frequency. This is a more practical way of analyzing sensors than just figuring out the total energy, because the spectral distribution of the noise is related to the frequency characteristics of the sensor dynamics.

Fortunately, this problem has been solved by Nyquist. Many other people since have given alternate derivations so it is easy to find in the literature. The key point is that the equivalent noise force squared is equal to 4 K_BT times the mechanical resistance times whatever increment of bandwidth you are talking about.

It is more productive in the long run to express the noise in power spectral density. Bring the δf over to the left side and call the right side the power spectral density of the force. Either way is fine.

Mechanical resistance is what? Force divided by velocity. Force times velocity is power. In general, there will be a "potential" quantity and a "flow" quantity involved. Their product is power; the potential quantity divided by the flow quantity gives an impedance of some sort. For electrical quantities, the potential is voltage and the flow is current. Electrical resistance is voltage over current: 4 K_BT times the electrical resistance gives you the equivalent voltage-squared noise. This is Johnson noise; everybody knows about this. Fewer people know about the mechanical version, and this can be just as influential in some sensors as Johnson noise.

At any rate, now you have a tool for figuring out how the energy is distributed over frequency. This is actually a low-frequency approximation. It is good as long as hf (Planck's constant times the frequency) is much less than K_BT. Fortunately, that is not much of a restriction. The frequency of interest has to be greater than 10^{13} Hz at room temperature before a more exact form is required. The simple form I've given implies the same power to infinite frequency but this is not possible. We do not have an infinite amount of energy spread over all frequencies. We won't get into any trouble with this low-frequency form but it is an approximation.

Now we can describe the distribution in frequency. Mechanical resistance is not necessarily constant with frequency. In fact, in most real systems it is not. But our expression does not require that R be constant. If R is a function of frequency, then the equivalent force-squared will also be, within a constant, the same function of frequency.

[Transparency I-14]

I am going to try to emphasize this next point several times. People lose sight of the fact that what you really want to do is detect a signal in an environment of noise. You can talk about just the noise of a system but it is easy to get jump to inappropriate conclusions.

I read a proposal just a few weeks ago in which a research group claimed to have developed an extremely low-noise sensor. That much was true. They had a very low-noise sensor. They neglected to give the signal responsivity, though. As it turned out, the signal responsivity was extremely low also. So the noise floor, even though it was quite low, was actually equivalent to a very large signal. (In other words, the sensor could only detect very large signals.)

It does not make sense to consider just noise or just signal in isolation if you are trying to build something that is actually going to work. It is very interesting to study noise by itself -- I find that study fascinating -- but you probably eventually want to make a useful sensor.

One way to avoid losing the relationship of signal and noise is to calculate the noise-equivalent signal. Take the calculated noise and figure out what signal would produce the same level at the sensor terminals.

For an accelerometer, an acceleration of the case creates a relative motion between the case and the proof mass. There is a mass, a spring, a damper, and a forcing function -- the noise-forcing function -- so that we do not forget that there is a mechanical component to the noise. The same is true for a pressure sensor. Don't just calculate the equivalent noise in terms of this case-to-proof-mass displacement; also calculate the equivalent noise in terms of what acceleration it is mimicking. That tells you something useful.

[Transparency I-15]

If you do that for the two simple systems in Transparency I-14, you get these two expressions. The noise force squared is equivalent to the mass times the noise acceleration quantity squared and that is equal to $4KbTR$ times df . If you do some rearranging and solve for the power spectral density of the acceleration, you obtain these expressions.

Once you have this expression, you can do some intelligent designing. If you know that this mechanical component is setting the noise floor on your system (and it may not), then you can design the system so that this noise is under control.

You might want to raise the Q of the system -- this is done by B&K in one of their low-noise microphones. Instead of blindly accepting the conventional design philosophy that the signal response of a sensor should be maximally flat over the whole band, they made the Q rather high. The principal resonance is just above the band of interest so the high Q produces a rise in response at the high end of the band. The net effect is to reduce the system noise floor. The response rise is compensated electrically.

You can increase the mass. That is a pretty obvious strategy. Depending on what else changes at the same time, the noise floor can be reduced. If you can increase the mass without changing the mechanical damping, you get quite a large increase. However, if the mass increase also increases the size of plates without increasing the spacing, the damping may increase also. The noise floor is still reduced but not by as large a factor.

MR. WAKELAND: You were saying that the response of the low-noise microphone is not flat across the band but that you can equalize it later in electronics. What about moving the resonance further up above where you are going to be operating? Does that work as well?

DR. GABRIELSON: That would flatten the response, but you will notice that as you move the resonance higher, the noise increases. What you would really like to do is lower the resonance, increase the mass, and increase the Q if you want a low-noise sensor. We have actually built one like that in conjunction with the Jet Propulsion Lab. We moved the resonance right into the middle of the band with a Q of about 10 or 12 and it works quite nicely.

Now, there are caveats, of course. If you have an array of sensors and the phase relationship is very important, then you will have to establish the relative phases some other way but, given that caveat, it is a nice strategy.

DR. FORTUNKO: There are some exceptions to this strategy. For example, if you are sensing a sinusoidal signal, it may be okay, but I have a situation where I am sensing an impulsive signal and my sensor has a high- Q resonance. The high- Q resonance picks up energy and having to cope with this high-level signal ultimately restricts the dynamic range, not the Johnson noise.

DR. GABRIELSON: Yes, if you have a Q of 3000, then you are probably going to waste a lot of dynamic range just handling that Q . However, if you have a Q of 10 so that the dynamic range is not excessive, you might not like the time-wave shapes that come out, but the spectrum results will be fine and the noise will be lower than a conventional design. There are limits to how far you can go, but the point I would like to make is that the classical design philosophy of critically damped resonance outside the top end of the band is not a smart design philosophy for a low-noise sensor.

DR. KEOLIAN: Even in the time domain, if you compensate electrically, won't the time signal come out exactly right?

DR. GABRIELSON: Yes. I was thinking of the case in which you did not compensate at all. Then the time-wave shape is distorted, but in the frequency domain you would be able to extract the part that you wanted. If you compensated electrically for the mechanical resonance, then the time-wave shape will be fine, too.

Again, you have to look at the system dynamic range. Will the system handle the levels at each stage? You can envision raising the Q of the mechanical structure so much that the mechanical structure cannot handle the range of signals that you want.

[Transparency I-17; skip 16]

Now I would like to work out a little problem. It is a very nice problem, because there are two seemingly disparate ways to get to the same answer. One of the nice things about solving problems in equilibrium thermodynamics is that it usually does not matter how you approach the problem. Different approaches lead to the same answer.

Consider a spherical source. The spherical source is radiating in the first breathing mode -- the entire surface is going out and coming in together. The pressure field away from such a source is described simply as long as we are not near any boundaries.

You can calculate the radial particle velocity from Newton's law: the relationship on the second line is $F = ma$ on a per volume basis. The mechanical radiation impedance in this case is the ratio of the force on the outer surface of the sphere to the radial particle velocity. The Nyquist relations that I wrote previously were written as $4 k_B T$ times R , the mechanical resistance. In general, it is $4 k_B T$ times the real part of the mechanical impedance.

The real part is the radiation resistance. If kr is very small, you get a simple expression for the noise-pressure spectral density -- a classic result. You often see this expression in papers on ambient noise in the ocean. They call this the thermal background noise in the ocean. The spectral density goes as frequency squared because the real part of the impedance goes as frequency squared.

In underwater systems, particularly, you start seeing this component at 20, 30, 40 kHz and above. There is, however, nothing to be gained by restricting the solution to very small spheres. If kr is replaced by the real part (inside the curly braces) of the radiation impedance -- that is, kr squared over the quantity $(1 + kr)$ squared -- then the noise expression is valid for arbitrary sized spheres.

Notice that now the thermal noise component increases as frequency squared until $kr = 1$ and then the noise becomes independent of frequency. This makes sense, because if the wavelengths of the noise fluctuations become smaller than the sensing transducer, they will average over the surface of it, because it is only the breathing mode that is excited.

MR. POESE: How is radiation a noise source?

DR. GABRIELSON: Any time there is a mechanism by which energy can leave the system, there is a fluctuation associated with it. Radiation appears to be a dissipation from the standpoint of the transducer.

MR. POESE: That was several slides ago, right?

DR. GABRIELSON: Right. It is the essence of the fluctuation-dissipation theorem. Radiation connects the sphere physically with the environment around it. If there is such a connection, then there is also a connection between the fluctuating environment at temperature T and that sphere.

MR. POESE: Molecules bouncing up against the sphere that you did not count on?

DR. GABRIELSON: Yes. Now, there is another way to treat this radiation problem that, on the surface, looks completely different.

[Transparency I-18]

It actually amounts to a derivation of the Nyquist theorem, but that won't be immediately obvious. Let's examine a hypothetical sensor, a point sensor for now, in the middle of a box with rigid walls. We will make it a cube to make the derivation easier.

There are acoustic modes associated with this box. If it is a rigid box, that means there are pressure maxima on the box faces. Because the system is symmetric, the modes have either a pressure maximum at the center or they have a pressure zero at the center. The ones with the pressure zero at the center do not contribute to the problem, because that is where our sensor is so the sensor cannot detect these modes. We need to consider only those modes with maxima at the walls and maxima at the center.

The mode shapes look like the cosine of the quantity (an integer times π times x/L). The shape functions have the same form in all three directions. The wave numbers, k , are simply an integer times π/L . Therefore the spacing between the k 's is π/L .

If you plot the allowable wavenumber points in k space, you build up a region that looks like a cubic atomic lattice. Each of these points is π/L from its nearest neighbor. What is the density of wavenumbers? The volume for a single mode is $(\pi/L)^3$ -- work that out for

yourselves. If cell volume is $(\pi/L)^3$ and there is one k point in that cell, then there are $(L/\pi)^3$ wavenumbers per unit volume.

Each of these modes can have kinetic energy and potential energy, so each mode reaches equilibrium with $k_B T$ joules of thermal energy: $1/2 k_B T$ for kinetic, $1/2 k_B T$ for potential. The k -space density of thermal energy is then $k_B T$ times $(L/\pi)^3$.

[Transparency I-19]

I do not really want the k -space density, I want the frequency-domain density. The rest of this problem involves figuring out how to translate from the δk domain to the δf domain. In other words, for any value of k , how many modes are in dk . Now, dk is a thin spherical shell in wavenumber space, so how many modes are in that shell? That is pretty easy; it is a geometry problem.

The energy in the shell, dk , is equal to the energy density times the volume of the shell. Divide by the spatial volume, and you get an expression for the volumetric energy density.

We also know that in terms of rms pressure, the volumetric energy density is pressure squared over density times the square of sound speed. Take these expressions, solve for pressure-squared, and you get exactly the result that we got before using mechanical radiation, a nice clean result.

[Transparency I-20]

It is interesting to think about how you would modify this derivation if the sphere were finite in size or if the point measurement were made on one of the rigid walls instead of in the center of the box. The changes are straightforward. You just have to count the proper modes that influence the problem. In the case of a point sensor on one wall, the sensor will respond to every mode in the direction perpendicular to the wall since each of these modes has a maximum on that wall. Only half of the cross modes have maxima at this point so you just have to be careful to count the proper set of modes. At any rate, if you do the counting and the algebra correctly, you get the right answer for point source in an infinite baffle. This slide shows the solution for a piston in a rigid baffle (dashed line) compared to the solution for a spherical transducer in free space (solid line).

I am going to skip *Transparencies I-21 and I-22* because they are straightforward explanations of the important noise components for a geophone. They illustrate a specific case of evaluating the spectral density of thermal energy for a particular kind of sensor. It does not involve any new concepts; it is just a step-by-step walk-through.

[Transparency I-23]

Here is a measurement of the noise floor of that geophone. The proof mass of this geophone is 32 gm, so the moving mass is not particularly small. If you are careful enough though, you can measure the noise floor, and here is the measurement. The large peaks are associated with suspension resonances in vibration isolation system I used to make the measurement.

If you predict the noise floor based on the thermal-noise mechanisms that I've described, you have to consider two components: the mechanical-thermal noise contribution and the electrical-thermal noise contribution. There is mechanical damping in the spring-mass system and there is electrical resistance in the coil of wire that is moving through the internal magnetic field. The smooth solid line is the prediction; the prediction has no adjustable parameters.

MR. POESE: How did you measure the noise?

DR. GABRIELSON: I built a special isolation chamber for studying self noise of sensitive acoustic and vibration sensors. The chamber has several stages of simple mass-spring mechanical isolation for reducing external vibrational noise. The chamber is a modification of the apparatus that Greenspan used to measure absorption in fluids. He hung a glass sphere on wires inside a vacuum vessel. My inner vessel is a stainless-steel cylinder hanging on springs. The springs are better for vibration isolation than wires. I pump a vacuum between the outer shell (a large, stainless steel bell jar) and the inner cylinder to isolate from external acoustic noise.

There is another way to do it: put two well-matched sensors in the chamber and connect them so that their signals are 180 degrees out of phase. You get a square-root-of-2 increase in the self noise and a large reduction in external signals if those external signals affect each sensor the same way. There are limits to how much cancellation you can achieve because the sensors have to be matched pretty well, but it is another option.

MR. POESE: And that cancels any real excitation?

DR. GABRIELSON: It cancels much of the real excitation, yes. It is not going to take a very high peak, though, and drop it right down into the background, but it can reduce interfering vibration or acoustic signals quite a bit.

MR. POESE: Is the high peak produced by HVAC equipment in the building or something like that?

DR. GABRIELSON: Indirectly. The system used for these measurements had a pretty flimsy tripod to support the inner chamber and there was a twisting mode in that tripod. It was a fairly high-Q resonance, so any external excitation was Q-amplified through the tripod. I've since rebuilt the inner structure to move these resonances much higher up in frequency.

DR. FORTUNKO: Why does the spectrum come down again beyond the resonance of the geophone?

DR. GABRIELSON: The two slides that I skipped give an explanation of the overall shape of the noise. The spectrum consists of two components: a flat (white) noise component that results from the Johnson noise associated with the electrical resistance of the geophone's coil; and a component associated with the mechanical damping. The second component is shaped by the mechanical resonance. The incoherent sum of the two components results in a small mechanical-noise peak around resonance rising out of the flat electrical-noise background.

[Transparency I-24]

Much work is being done in fiber optic sensors. Self noise is important in any sensor design and the aspect of noise is particularly fascinating in fiber sensors because there are so many different varieties of noise mechanisms.

Almost every noise mechanism you can think of can be found in a fiber optic system. That does not mean they are inherently bad; it just means you have to be very careful if you are designing low-noise fiber sensors.

To illustrate, here is silly fiber sensor. It consists of an optical fiber clamped at one end with a mass hanging on the other end. You would construct an interferometer with this leg and a reference leg.

The fiber and the mass form a mass-spring system because the fiber has some elasticity. If you accelerate the support, you would be able to measure the acceleration by the optical phase shift produced by the change in length of the suspension fiber. The first thing you know is that there is an ordinary mechanical-thermal noise level set by the mechanical damping of this system.

If you pulled this mass down, not far enough to break the fiber but just far enough to stretch it a little bit, and let it go, the mass would oscillate up and down and you would be able to measure the Q. If you know the mass, you would know everything you need to know to find the mechanical-thermal noise level.

However, the mechanical damping is not the only source of dissipation in this system. The fiber itself has loss. One of fiber's advantages is its low optical loss, but it still has loss. If it has loss, then there must be a fluctuation associated with that loss.

In fact, there are at least two loss mechanisms that can be identified in theory, although it is very hard to separate them experimentally. One is scattering from thermally induced density (or index of refraction) fluctuations in the fiber. Another is scattering because of nonuniformities in the fiber, which are either impurities or density fluctuations that were frozen in when the fiber was cooled as it was made. In principle, measurements made at a number of different temperatures could distinguish between these two but I don't know if this measurement has ever been made.

DR. KEOLIAN: Why is the second, the frozen optical scattering, a dissipation and not a reactance? What makes it random? How do these frozen inhomogeneities cause energy to leave the system?

DR. FORTUNKO: Eventually the scattered light gets absorbed. It basically changes the energy from moving in one direction into some other direction so it takes energy out of the system that way.

DR. MIGLIORI: If the fluctuations change the angle that the light hits the fiber core wall so that the light can escape, then they look lossy. Otherwise they do not. This is the same thing in resonant ultrasound: you can put an array of scatterers in and it does not change the Q, because there is no path for the sound to get out. But here you can scatter and every once in a while one of those fluctuations will knock a photon out of the fiber.

DR. FORTUNKO: Also, the fiber can change dimension with time because of creep and that produces another kind of noise.

DR. GABRIELSON: Yes. There are a number of very interesting systems that have noise associated with creep. Sensors constructed with piezoelectric ceramics can produce noise directly associated with creep. I would guess that this is probably a much more serious problem than it is in fiber, because, in the piezoelectric material, the mechanism couples into the domain shifts and randomly creates quite a bit of charge. Creep noise or stress-induced noise in ferroelectrics I have measured and I would not be surprised if there were something similar when an optical fiber has some static prestress.

[Transparency I-25]

The expressions on this slide give you a rough idea of the mechanisms. There are two terms inside the brackets: the change in index of refraction with temperature and the change in

length with temperature. Temperature fluctuations in the fiber will produce signal (optical phase) fluctuations through both effects.

For typical single-mode fibers there is a frequency dependence to noise and it turns out that the equilibrium phase fluctuations from these mechanisms are important from very low frequency to 20 or 30 kHz or so. If the signals of interest are in this frequency range, then this source of noise is important.

According to this theory, the location of the frequency break point is determined by the size relationship between the thermal penetration depth, which is the distance through which heat can diffuse in an acoustic cycle, and the mode-field radius of this fiber. This theory was developed for a fiber in an infinite insulating cladding, but the essential details do not change with finite cladding as long as the cladding is much thicker than the mode-field radius.

This is a very interesting relationship between the mode-field radius and the thermal penetration depth. At frequencies above the break point, the thermal penetration depth is small with respect to the optical cross-section, so all the fluctuations are effectively being averaged over by the mode field and they don't contribute much. It is analogous to the problem of averaging over the surface of a finite sphere that causes the noise to make a transition from increasing as frequency squared to being constant with frequency.

MR. WAKELAND: Can you say again what the mode-field radius is?

DR. GABRIELSON: A fiber guides with some kind of index-of-refraction change between the core and the cladding. It might be a step change, it might be a graded change. In each case, you would find that the optical energy-distribution profile would be different; however, in each case, the optical energy is confined. The mode-field radius is a measure of the width of that energy-distribution profile in the fiber; its value is not too far different from the physical core radius.

DR. MIGLIORI: Let me understand this. If you are at frequencies above which the thermal penetration depth equals the mode-field radius, then the noise goes away?

DR. GABRIELSON: It starts being averaged out across the diameter of the fiber so the noise power drops with increasing frequency.

DR. GARRETT: Why do the fluctuations average?

DR. GABRIELSON: The optical field is coherent over the entire mode field diameter in the fiber. A fluctuation that extends over the entire diameter fully affects the optical field. The thermal penetration depth is related to the coherence distance of thermal fluctuations. Very crudely, you could imagine that there is one independent fluctuation in each patch of diameter equal to the thermal penetration depth. If the thermal penetration depth is much smaller than the mode field radius, then there would be many independent fluctuations across the mode field cross-section. Each fluctuation is acting on the same coherent optical wave front; each fluctuation is independent, so the net effect on the optical field is an average over those fluctuations.

A more rigorous explanation would involve the coherence distance of the fluctuations. Again speaking roughly, the fluctuation is coherent over the scale of a thermal penetration depth. If you remember our discussion of the noise experienced by a spherical "transducer" when we considered the case of ka larger than one we found that the noise stopped increasing as frequency squared and became constant. An alternate way of looking at this problem is that, for $ka > 1$, the

coherence distance of the fluctuations is smaller than the sphere and so they start averaging out over the sphere's surface.

[Transparency I-26]

The last thing I am going to talk about in this section is a very clever theorem to save you some time in calculating signal-to-noise ratio in complicated systems. This theorem says that for a linear system the signal-to-noise ratio at the output of that system is independent of the load impedance on the system output. You can often greatly simplify the noise calculation by shorting the output. You can do that either by driving a resistor and letting the resistor go to zero or you can drive an actual short and calculate the current. If you do the calculation for both signal and noise voltage or current (an external signal source and internal Johnson noise for instance), then form the signal-to-noise ratio, you'll get the same answer as you would for the original system with non-zero output load.

Feedback systems are usually tedious to analyze. By using this theorem, you can often find the answer much more quickly. In the example shown, shorting the output breaks the feedback loop and the resulting calculation is easy to do. The calculation does not tell you anything about the noise or the signal individually; you would have to do those calculations using the proper output impedance since both the signal and noise transfer functions depend on the output impedance. The ratio of signal to noise, which is often the crucial quantity, is not a function of the load impedance.

[Transparency I-27]

MR. POESE: When you do the noise calculation do you take out all the noise sources except one and figure it out and then put another noise source in and figure it out and add them up?

DR. GABRIELSON: That is exactly what you do and then you ask yourself the question, are the noise sources in the system coherent to any degree, or incoherent. In many cases, they are incoherent and you can just add them as RMS values, in other words, square them, add them up, and take the square root.

There are important cases where they are not incoherent. The fortunate thing that saves the casual worker is that the error is relatively small in most cases by assuming incoherence. However, it is not correct to say that all noise sources are incoherent; some are caused by the same mechanism but appear in two different ways. If two of the noise sources in a system are perfectly correlated, what is the maximum error that you would make? If the two sources produced equal levels, then the RMS sum would produce a total 3 dB higher than either. If the two sources summed coherently, the total would be 6 dB higher than either. So the maximum error is 3 dB (or a factor of 1.4 in amplitude). Usually though, the two components do not produce the same amplitude and then the error is smaller. If one source is half the amplitude of the other, the maximum error would be 1.6 dB (or a factor of 1.2 in amplitude). If a third source produces an amplitude five times larger than either of the other two, then the error would be, in most cases, immeasurable.

The message is to understand the phenomena but don't spend a lot of effort on effects that are swamped by something more consequential. Understanding the effects of all of the noise sources is necessary to sound design: find the dominant noise mechanism and work to reduce it (all the while considering the noise-equivalent signal so you don't reduce the signal response

faster than you've reduced the noise response); once that mechanism is somewhat below another one, it doesn't pay to continue to reduce it -- move on the new dominant mechanism.

Part II: Shot Noise and Non-Equilibrium Noise

[Transparency II-1, 2]

I'd like to pose a another question to introduce the idea of shot noise. What happens in this circuit? If you measure the noise voltage across the resistor with no current flowing through the resistor, you will get a number that is very close to the Johnson noise -- $k_B T$ times R -- if you have done the measurement carefully; if the amplifier noise has not swamped the resistor noise.

I used to say that this measurement was easy to do but Professor Garrett has convinced me to stop saying that. I admit that it is easy only after you have done it 10 or 20 times, but it is definitely a straightforward experiment; it is not something beyond the capabilities of anyone comfortable with simple circuits and a spectrum analyzer.

What happens if you put a current through that resistor? How does the noise change when the current is going through? Let's assume that the current is small enough and the resistor is well enough connected (thermally) to the environment so that the resistor doesn't change significantly in temperature.

An electrical current is made up of discrete particles, electrons in the classical view, so there ought to be noise associated with random arrivals of those discrete particles -- something known as shot noise.

[Transparency II-3]

Before I tell you what the answer to my question is, I am going to introduce shot noise. If you have a system that is characterized by the flow of discrete particles or discrete wave packets, then you can identify an energy associated with each isolated lump. If the particles or wave packets are striking your detector in some random and independent way, noise is produced, and it is straightforward to calculate how much noise.

There are several ways of doing the calculation. Here is one of them. Let's say that we have an ordinary electrical current that is made up of packets each carrying a charge, q . They are sufficiently concentrated in time that we cannot resolve their features with our measuring system, so we are going to call them delta functions. In the ideal model, they have the same height, they all carry the same charge, but they arrive at random times, so you can represent it as a sum of randomly occurring delta functions.

If you expand the current as a Fourier series, that is, evaluate the coefficients for the series, you can calculate any particular component of that Fourier series. Then you take the mean square value of the "a" and "b" coefficients, which is $1/2$ the sum of the squares of the two coefficients.

What we have calculated is the amount of energy in the one increment of frequency that this particular component represents. From the sampling theorem we know how wide that increment is. It is $1/T$, where T is the averaging time. What remains is to explicitly evaluate the time averages of the squares of the "a" and "b" coefficients.

[Transparency II-4]

So we square the expressions for "a" and "b", add, then time average. This gives you two classes of products. It gives you products of like functions with the same arguments, that is, the

diagonal elements, and it gives you products of functions with different arguments: the off-diagonal elements.

There are some cases in which the time averages of the cross-product terms go to zero, but it is important to realize that this is not always true. It depends on the independence of these quantities. Consider, for example, what happens if events happen with regularity; they happen every microsecond with 1 microsecond between them, not a bit more, not a bit less. The cross-product terms do not go to zero in this case. If the events are independent -- if you can't predict anything about the next event even if you know everything about the previous events -- then the cross-product terms do average to zero. In that case, you get a simple expression for the time average of the square of the "a" and "b" coefficients. The mean current is N times q/T . The bandwidth is $1/T$. If you take this expression and insert it into the expression on the previous slide for i -squared, you get the following mean-square value of noise: 2 times q times the mean current, which is the classic shot noise expression.

This result applies to processes consisting of events that are impulse-like and independent. We can relax the condition that the events be impulses without too much difficulty. Things are more complicated if one event shows dependence on other events. You have to evaluate the cross-product terms. At any rate, now you know the form of shot noise expressions and the conditions required.

[Transparency II-5]

This circuit is the first circuit that I used for this particular measurement. It is a simple low-noise pre-amplifier design. The circuit front-end is a JFET and, as for any transistor, the JFET requires a DC path from the gate to ground. Normally, this path would be a single resistor. If instead of using a single resistor, you make a bridge-like structure of matched resistors, then you can put pretty much whatever you want in the middle of the bridge without affecting the electrical characteristics that the JFET gate sees.

If you put 4 one-mega-ohm resistors in an arrangement like this, the equivalent resistance from the gate terminal of the bridge to ground is 1 mega-ohm, whether the switch is closed or not. If this battery is not here and I just short across these terminals, it is still 1 mega-ohm, I have not changed anything. By closing this switch, though, I can put a sizable current through these legs and, if that current produces shot noise, I can easily put enough current in to overwhelm the Johnson noise.

[Transparency II-6]

Here is an actual measurement from this circuit. The B curve is with the switch open, that is, with no current flowing through that combination of resistors on the input to the amplifier. The measured level agrees very well with the Johnson noise prediction for 1 mega-ohm.

Close the switch. Curve A is the resulting noise. The prediction for shot noise based on the DC current that was produced by the battery is shown and that prediction is much higher than the measurement. In fact, except for the low-frequency end, the measured level is not discernibly different from the measurement without any current.

At the low-frequency end there is some $1/f$ noise, which is more complicated than any of the mechanisms we have talked about so far. The cause is probably related to the "wiping" contacts between the resistor leads and the sockets in the prototyping board that I used to assemble the circuit. The character and level of the $1/f$ noise does not change much with the type

of resistor (carbon, carbon film, or metal film) and this would not be the case if it were an effect internal to the resistor: carbon resistors have much higher $1/f$ noise than metal film resistors. The contact noise is overwhelming the resistor $1/f$ noise. The important point here though is that, except for the $1/f$ noise, the noise does not increase as the shot-noise expression predicts. And, it is not a subtle effect; shot noise is clearly NOT present.

The bridge arrangement, by the way, not only provides a means for introducing the current without upsetting the operating point of the transistor. The bridge also rejects noise from the battery. The current from the battery goes through both legs of the bridge equally (within the limits of matching resistors in the bridge). If there is noise produced by the battery, the noise will show up equally in both legs and cancel at the measurement terminals. You ought to think about that to be sure you understand it. The naive way to do this experiment is to force a current through a single resistor hung on the transistor's gate. There's no guarantee what you'd be measuring in that case.

PARTICIPANT: Why didn't the shot noise show up?

DR. GABRIELSON: I am going to talk about that shortly. I wanted to set the stage first. Before I answer that question, I want to work through another problem. I promise I'll come back to it though.

[Transparency II-7]

Let's try to analyze the noise that results from molecular collisions. This is a bit more applicable to acoustics than noise in resistors. Consider a disk; molecules are hitting this disk. As we did with thermal noise associated with radiation resistance, we are going to do the calculation two ways to see if we can get the same answer.

The force on the disk is equal to the sum over all molecules of the momentum transferred in the collisions. I can write that as molecular flux (molecules per second per square meter) times momentum change per collision: that gives the force on the disk.

In this example, the disk is free to move (like a microphone diaphragm, for example). It can move with some velocity, \dot{x} . The molecules have some statistically averaged velocity too. The molecular flux is the number density, n , of molecules times the area of the disk times the difference in these two speeds. (The factor of one-half is introduced because only half of the molecules are moving toward the disk, the other half are moving away.)

The momentum change is 2 times the mass of a molecule times the relative velocity between the disk and the molecule because the molecule is changing direction, it is not just sticking on to the disk. Expand the force expression and you get three interesting terms. The first term is the static pressure on the disk.

The second term is a coefficient multiplying \dot{x} , so it must be a damping term. I am going to call it R_{mech} for mechanical resistance. The third term is proportional to the square of \dot{x} . I am going to restrict the motion of the disk so that \dot{x} is always much smaller than the average thermal speed of the molecule. This is an entirely reasonable assumption in most cases.

[Transparency II-8]

Since we now have the mechanical resistance, we should be able to calculate the fluctuation pressure. The fluctuation pressure squared, which is the equivalent noise force squared divided by the disk area squared or $4k_B T$ times R_{mech} divided by the area squared. From kinetic theory, I know that the static pressure is n times $k_B T$. For a Boltzmann distribution

of velocities, the average velocity is equal to two times the average of any one of the x-, y-, or z-component velocities. This leads to the expression shown for the fluctuation pressure squared.

The pressure squared is equal to 2 times the quantity, $2m$ times the average velocity, which is something like the amount of momentum transferred by these molecules, times the static pressure divided by the disk area. This has something like the form of a shot noise-expression.

[Transparency II-9]

If you write the electrical shot noise differently, though, it looks exactly like this pressure noise expression. We wrote the electrical version as it usually is written: the current noise squared is equal to 2 times the charge times the average current. A more general way to write the shot noise expression is as follows: the mean square fluctuation in flux density is equal to 2 times the quantity being carried, whatever the relevant quantity is, times the average flux density times the bandwidth divided by the area.

Electric-charge flux density is J , so the flux-density fluctuation j -squared is equal to 2 times q times the average flux density, J -zero divided by the area. If you convert back into currents, you get the usual shot-noise expression. Photon-flux density is optical intensity. The fluctuating part squared is equal to 2 times hf (hf is the energy, the quantity being carried) times the mean intensity divided by the area. For momentum-flux density or pressure, we get the same expression as before.

If we really believed this shot-noise formalism, we could have just written the expression for fluctuation pressure, but it was helpful to go through the process of constructing the expression so I can tell you why it is wrong.

The expression that we just derived says that the pressure fluctuation noise power is proportional to static pressure.

MR. POESE: Does that mean that the dissipation -- we are back to the dissipation-fluctuation theorem -- opens the door to this shot noise and is based on the fact that the microphone diaphragm is allowed to interact with the atmosphere?

DR. GABRIELSON: Yes, with the caveat that I just did the whole derivation that you have seen on the last two slides incorrectly.

[Transparency II-10]

I am going to do the derivation another way now, because I ought to be able to do it in a different way and get the same answer. Equilibrium thermal fluctuations in force are given by our old friend the Nyquist relation. The frictional force on a disk in a gas is a well-studied problem first solved many years ago.

The mechanical resistance is 16 times the viscosity times the radius of the disk for broad-side motion of the disk. Plug that into Nyquist's relation and what do you get? The fluctuation noise power is virtually independent of static pressure. Viscosity is a weak function of pressure, much weaker than linear with pressure. This expression is radically different from the previous one. So something is wrong.

[Transparency II-11]

The shot-noise expressions all require that the collisions be independent and there is nothing at all independent about the collisions of gas molecules at normal pressures with a disk. It may not be obvious at first, but consider the situation if the disk is moving, and there are lots of molecules, not just one. The drawing I made was misleading. A single molecule does not just

start out somewhere, hit the disk, and rebound back out. The region around the disk is crowded with molecules. In a traffic jam, your car is surrounded by many, many other cars; the motion of any one car depends greatly on the motion (or lack thereof) of the surrounding cars. Same idea with molecules except that the mechanism of interaction is collision rather than near-collision (and horn honking).

If a molecule hits the disk, it's motion is affected by the motion of the disk. It rebounds from the surface of the disk and immediately interacts with many other molecules transferring its new momentum value out into the bulk of the flow. So most of the molecules that are inbound to the disk have had their velocities modified before reaching the disk by the rebounding molecules. What I assumed in that first derivation was that the motion of the disk had no influence on the statistics of the molecules near the disk and that is very wrong at atmospheric pressure.

That is why the derivation produced nonsense. It was absolutely wrong to assume that the collisions were independent. They are not. They are highly dependent. A molecule on its way in interacts with many outbound molecules before it hits the disk and then it, in turn, interacts with many molecules on the way out.

DR. MAYNARD: This is the same mistake that is made in a lot of textbooks when they describe Brownian motion. A lot of books claim that Brownian motion shows that gases or fluids are made up of atoms. It does not show that; it just shows that the gas or fluid has viscosity.

DR. MIGLIORI: If you looked at a box and you looked at all the atoms going left and all the atoms going right, you would have a certain distribution if it were in thermal equilibrium. Suppose one of the walls is moving: the molecules going left have a different velocity than the molecules going right. As soon as they interact, they equilibrate.

DR. GABRIELSON: Right. If they do not interact, they do not equilibrate, and the first derivation would then be fine. In fact, this happens if the gas pressure is so low that the mean free path is much larger than the disk. What I really did was a derivation for the case in which the mean free path is much larger than the disk.

What does the fluctuation pressure as a function of static pressure look like, then? In the free-molecular region, the pressure fluctuation does increase with the ambient pressure, but at the point where the mean free path is equal (more or less) to the radius, we move into another region, that of normal viscous flow, where the fluctuations are virtually independent of static pressure.

The region between free-molecular flow and normal viscous flow is more complicated than you might suspect. This region is usually called slip flow. Here, the flow behaves as if it is a continuum viscous flow but with slipping at the boundary. As a consequence, the transition region is very large. There is a 10% effect on the flow when the mean free path is just one sixtieth of the disk radius. This is a lot larger than you would guess. This transition region is very interesting. We certainly know how to handle the free-molecular region, and we know how to handle the viscous-flow regime, but the transition region is tougher.

DR. GARRETT: That drawing is a little misleading, because when you look at it you think of an RC time constant. But it is really not so simple as that. The effects of this transition region are very broad.

DR. GABRIELSON: Right.

[Transparency II-12]

How does all this relate to our experiment in electron conduction through resistors? Before I explain why the measurement we made was surprising, I'm going to give a brief introduction to noise in different types of conductors. In a metallic conductor, there is a band structure that is related to the interaction of the electron wave packets with the periodic lattice of atoms in the metal. I presume that Dr. Maynard could set up a simulation for a wave packet in a periodic lattice and show the origins of this band structure. At any rate, a metal is characterized by a conduction band that is not full but does have plenty of carriers.

The Fermi level is, roughly speaking, the average height of the energy "pile" produced by the conduction electrons. I have drawn them as if they are just piled-up eggs; that is not really right; I'll show you a better representation later.

The carriers in a metal are very highly correlated. If you think about it, you might suspect this. The electrons are charged particles (in the classical view). If they are close together, then their electromagnetic fields will interact and produce dependency. The noise is independent of the flow rate because of that dependency or correlation. The equivalent current noise is the electrical-current form of the Johnson-noise expression -- $4 k_B T$ over R -- not the shot-noise expression.

[Transparency II-13]

In a semiconductor (an intrinsic semiconductor) every carrier gets into the conduction band by thermal excitation. At absolute zero temperature there are no carriers in the conduction band and the material behaves as an insulator. As the temperature increases, carriers are kicked up into the conduction band. If the concentration of carriers in the conduction band is sparse enough, the carriers act as if they are independent, and the current-noise fluctuations are well described by a shot-noise expression.

If I had replaced the 1 mega-ohm resistors (which were metal film resistors) by very lightly doped semiconductor resistors, I would have seen the full shot effect in that circuit.

DR. GARRETT: Do carbon-composition resistors qualify as semiconductor resistors?

DR. GABRIELSON: No. Their carrier concentration is too high. I've done the experiment with carbon-composition resistors and you get the same result, except that the $1/f$ noise is higher. Above the $1/f$ region, however, the level is Johnson noise.

DR. MIGLIORI: You should try this with a 10 mega-ohm thin-film resistor sometime, because the conduction starts to become two-dimensional.

DR. GABRIELSON: Good idea.

So far I've given a description of the difference between a metal and a semiconductor without any justification. Now I'd like to dig a little deeper.

When I was solving the problem of the noise associated with radiation from a sphere, I used a mode sum of waves in a box and I let the box dimensions go to infinity. You'll remember that I found evenly spaced k values in three dimensions and sketched the distribution of k values as a cubic lattice in three-dimensional k -space.

You can do the same thing for conduction electrons. Electrons are fermions -- I can use that word, because Dr. Maynard has already explained what it means. Two fermions cannot occupy the same spin state. This is often stated differently. At most, only two electrons can share a given state: one will have an "up" spin and the other will have a "down" spin. So instead

all the electrons being at one energy level, the allowable k values (states) are filled from lowest to highest with the electrons having increasingly higher energies.

[Transparency II-14]

If I just draw the two-dimensional representation of the cubic lattice of allowable k values, I create a grid of evenly spaced points. In a crystalline metal the spacings and overall shape of the distribution can be pretty bizarre, because the crystal is not isotropic. Usually though a piece of wire is not a single crystal but a collection of many crystalline grains with more or less random orientation. Normally, the evenly spaced cubic k lattice is an acceptable simplification under these circumstances.

If I put a solid disk centered on top of the transparency, the shadow on the screen shows what a metal's k -state lattice looks like at zero temperature. The states are filled from the k -space origin outward until you have used up all the available conduction electrons. A two-dimensional slice at zero temperature looks like a solid disk with a sharply defined edge. (Now I'm using a grid of k values that has very small spacing because I need to account for very many electrons; therefore, the filled portion appears to be solid rather than a grid.) In three dimensions, the occupation would look like a sphere with a sharp boundary.

For copper, the energy associated with the outside of the occupation disk is 7 eV, which is the average Fermi level for copper. That is huge compared to $k_B T$ at room temperature, which is about 0.025 eV. If there is an electron packet somewhere in the interior of the disk, thermal excitation cannot move it because there are no empty states nearby. The only empty states are near the outer ones, so when the temperature goes up, only the electrons near the boundary are affected. I can show the effect if I raise the disk above the transparency a few inches so that the projected image is blurry around the edge. [Applause]

At room temperature, the interior is unchanged but outer edge of the occupation sphere is fuzzy. These outer electrons are the ones responsible for the observable behavior of the metal, whether it is the DC current flow or the noise or the specific heat.

There are really 2 viewpoints. You could consider that every electron contributes to flow, for example. If I put a voltage across the conductor, the occupation disk moves a little away from the origin; it moves very little even for a large voltage. To displace the center of the occupation disk one radius (7 eV) away from the origin, I would have to put so much current through the conductor that it would vaporize. If I put a normal voltage across the conductor, the disk slides just a bit (to the right in my coordinates).

As you can see, most of the momentum states have corresponding states that cancel out: a point in the first quadrant is cancelled by a point in the fourth quadrant and so forth. When the disk is perfectly centered and at zero degrees absolute, they all cancel out. When the disk is centered but at room temperature, the states in the fuzzy outer region cancel on average so that there is no DC current. But the time-varying imperfections in the cancellation produce observable noise. When the disk is off center because of an applied voltage, there is some net momentum left after cancellation. So you can either consider that all the electrons are active but most of them cancel out or you can consider that only outermost electrons -- the ones that don't cancel -- contribute to either the DC flow or the noise.

The point in terms of noise is that the fluctuating part is pushed way out onto this thin, outer shell and moving the disk slightly has very little effect on that shell. A bit of asymmetry is

introduced but the effect on noise is negligible. There is a, perhaps large, DC flow but the noise does not change. I will show you analytically in a little bit why that is true, but just from the figure you can see that it might be true. It is very much like having two counterbalancing noise currents; they each shift by only a tiny bit when you apply a voltage, so there isn't much change.

[Transparency II-15]

What happens in a semiconductor? A semiconductor at zero degrees Kelvin has no carriers in the conduction band at all. There is no occupation disk. If you raise the temperature, though, the center region becomes faintly fuzzy as it does on the transparency when I hold a very small disk a few inches above the surface. A small, out-of-focus blob appears. The whole region is fuzzy and much smaller even than I'm showing here. There is no large central core of very well-defined energy momentum states. Everything is fuzzy. When you apply a voltage now, the fuzzy disk slides over a little bit. In this case, there is a significant difference between the shifted and unshifted cases and the character of the noise changes. Now you get shot noise.

I've been talking about an intrinsic semiconductor or, at most, a lightly doped semiconductor. If the semiconductor is strongly doped, then there are enough conductors to have some correlation. Most interconnection paths in integrated circuits are semiconductors but they are highly doped so they act like metals from the standpoint noise. Otherwise, it would be very difficult to produce low-noise amplifier chips.

MR. WAKELAND: Would you say again what this picture represents?

DR. GABRIELSON: You are looking at a map of the modes of a box. If you calculate the modes for a box, whether they are acoustic or electromagnetic, they each have a wave number and the higher modes have higher wave numbers, but they are discrete, so you can do a point-by-point mapping into wavenumber space. The allowable modes form the occupation region. The background in the picture is the grid of wavenumbers.

MR. WAKELAND: When you put the disk on the transparency when you were talking about conductors, what did the disk represent?

DR. GABRIELSON: It represented how far out the allowable modes filled the diagram. Another way to say that is that it represented the filling from the origin outward of allowable states by the conduction electrons. The electrons (or the equivalent wave packets) fill in from the origin outward until there are no more carriers to place on the diagram. Of course, the grid is a lot finer than I have drawn here.

MR. WAKELAND: We are talking now only about the conduction band?

DR. GABRIELSON: Yes.

MR. WAKELAND: And that is why you are saying that for a semiconductor at a low temperature, there's nothing on the diagram?

DR. GABRIELSON: Right.

[Transparency II-16]

To draw this another way for the conductor, slice through the occupation disk from the origin out and just plot the likelihood of occupancy, which is one if the state is always occupied, zero if it is never occupied, and some level in between if the state may or may not be occupied. The result is a curve that is given by the distribution function that Dr. Maynard described. The occupation is one all the way out to the near the Fermi level. If the system is at zero temperature, it drops abruptly down to zero. If the system is at room temperature, then there is a small

transition region, and that is where all the interesting stuff is happening. That is the fuzzy area in the occupation disk. The curve labeled "df" gives the derivative of the occupation curve and it shows where occupancy changes will occur. It's zero everywhere but in the thin, fuzzy shell.

[Transparency II-17]

In a semiconductor you notice two differences: now the curves "f" and "df" are identical. Everywhere there is occupation, and there is change in occupation. Also, notice that the absolute occupation level is a lot lower, too; the occupation probability is much less than one everywhere.

You can do a nice measurement of this semiconductor shot noise in a bipolar junction transistor because the current noise can be made to overwhelm all other sources of noise. I did this because, at PASS two years ago, Professor Garrett mentioned that, in op amps, if you measure the leakage current, you cannot predict the noise current to better than a factor of 2 or 3 or so by the shot noise equation. The current noise is often 2 or 3 times larger than you would calculate from the square root of two times q times the leakage current.

I still don't know why this factor of 2 or 3 is commonly found, but I was able to make a decent measurement on a system that I think I understand fairly well: a bipolar junction transistor. The leakage current is fairly large and, if a reasonably large base resistor is used, the output noise is dominated by the voltage produced by the input current noise flowing in to that resistor. Measure the leakage current into the base and also measure the output noise voltage. Convert the output noise voltage into equivalent input current noise (divide by the voltage gain, then divide by the base resistance). The shot noise calculation based on the input leakage current agrees with the measured input current noise within a few percent. So, in that system, the shot-noise model is a good one.

[Transparency II-18]

Earlier, I mentioned the idea of having two currents, one flowing in one direction and one flowing in the opposite direction. You normally can't measure the currents individually, but it's a useful model of the real behavior. The external current that you do measure is the difference between the two opposing flows.

If there are two such opposing currents, the measurable current is the difference between the two. The equivalent noise current, if the two flows are uncorrelated, is the square root of the sum of the squares of the two noise currents.

There are a number of systems in which this model of two opposing currents is useful. Tunneling through a potential barrier is one. Flow across a P-N junction in a diode is another. In both of these cases, the relationship between the reverse current and the forward current is given by a Boltzmann-law relationship: the reverse current is equal to the forward current times an exponential. The exponent is minus qV divided by KbT . This form -- "e" to power formed by the ratio of some energy (qV in this case) to the thermal energy (KbT) is common to many problems in systems with distributions of energy. Having this relationship between the forward and reverse currents is all that is required to find an expression for the equivalent noise current in terms of the measurable ("DC") current.

There are two interesting regions. If the voltage across this system is relatively small, that is, qV is much smaller than KbT , the expression reduces to the expression for Johnson noise. It does not matter what the mechanism is; if you are not pumping energy into it, you are only going to get equilibrium noise out of it.

If qV , though, is much greater than $k_B T$, the expression becomes the classical shot noise expression. This is very nice. You get the right answer for equilibrium and if you pump energy into it and push it away from equilibrium, you get shot noise.

[Transparency II-19]

This slide contains a figure from a very recent paper (1997). It is one of two papers (by different groups) that were published recently on measurements supporting the fractional quantum Hall effect. I am not going to attempt to explain what the fractional quantum Hall effect is: I don't understand it myself. But both groups wanted to establish that in a particular situation electrons were behaving as if they were particles with charge $q/3$ instead of particles with charge q .

What caught my attention was that both groups demonstrated this by measuring the noise of the system. Their system is also well modeled by two opposing currents. Here are the measurements from one group. The asymptote at zero current is the Johnson noise region. The asymptote at high current matches the shot-noise model as long as you use $q/3$ instead of q in the model. The dotted curve that climbs sharply with increasing current is the shot-noise prediction for a charge per carrier of q . This is not a subtle difference.

[Transparencies II-20 to II-23]

There are a few more pages in this shot noise/non-equilibrium noise section that I am not going to talk about. They are pretty self-explanatory and the section on $1/f$ noise introduces a topic in which there's a great deal of work still being done. You can read more than you would ever want to know about $1/f$ noise, and I suspect that some of it is even accurate and useful.

I used to read all the $1/f$ noise literature I could get my hands on, but I eventually gave up. I burned out. It's hard to separate the wheat from the chaff. It is a fascinating subject, though. If you want to start a study of some aspect of noise, it is an interesting one.

Part III: Sensor Calibration

I would like to spend the last hour talking about some techniques for sensor calibrations. I am going to talk about two techniques for absolute calibration: one is reciprocity and one is an optical interferometric technique. I'm going to concentrate on techniques that are practical to do in the laboratory and also capable of reasonable precision. (Note: In this section, complex quantities will be used. The symbol, j , will be used as the indicator of imaginary quantities. Where the symbol, i , appears, it indicates current.)

[Transparency III-1]

Comparison calibration is used more frequently than either of these techniques but comparison calibration is not nearly as interesting to discuss and is, of course, not capable of absolute calibration. Furthermore, comparison calibration is not always practical. For example, your experiment may involve such an extreme environment that you cannot calibrate the sensor outside the apparatus, then put it in the apparatus and make your measurement. The environment may cause a significant change in the sensor's response. In this case, you may want to make an in-situ calibration every time you start a measurement. Often the apparatus can be designed so that a reciprocity calibration can be done without disturbing the environmental conditions or the apparatus.

[Transparency III-2]

Consider a system with two ports. Each port has two terminals. This could represent an electromechanical transducer of some type. A voltage may appear on one set of terminals when a current is forced into the other set. If the device is a transducer, a velocity might be the forcing "flow" instead of current. The two-port system can have many combinations of physical quantities at the ports depending on the circuit or transducer that it represents.

If a transducer (or a circuit or a transmission medium) is linear and it is reversible -- that is, it functions in both directions -- there is a pretty good chance it will also be reciprocal. The definition of reciprocity is more specific, though. To understand the definition, we will write equations that characterize the transfer through this system. In the example shown, the "potential" quantities are voltage and force, and the "flow" quantities are current and velocity. The voltage at the electrical terminals can be written as a function of the two flow quantities. Similarly, the force at the mechanical terminals can be written as a function of those same flow quantities (all the time assuming that the system is linear). When you have written those equations, if you discover that B equals C , the system is reciprocal.

One of the most important consequences of reciprocity is that the ratio of voltage to velocity for zero current (i.e., the open-circuit voltage) is equal to the ratio of the force to current for zero velocity.

[Transparency III-3]

There are cases where you might find that B equals $-C$. Such cases are often called anti-reciprocal and there is a whole body of literature that talks about antireciprocal devices. But, because the behavior is governed by two linear algebraic equations, you can always transform the equations into a form in which B equals C . This process interchanges the potential and flow quantities at one of the ports. One of the most common transformations is done for moving-coil transducers (loudspeakers, for example). This leads to velocity being the potential quantity and force being the flow quantity at the mechanical port.

DR. KEOLIAN: Can you always transform an anti-reciprocal physical system into a reciprocal system by changing variables on the ports?

DR. GABRIELSON: No. You can always transform a SINGLE two-port block from anti-reciprocal to reciprocal. A physical system might require more than one two-port block for representation (or it might not be representable at all with two-port blocks). In that case, the ports of different blocks that are connected must be compatible: if force is potential for a port of one block, force must be potential for the port of another block if those two ports are connected. Consider a moving-coil transducer ("anti-reciprocal") that is connected in series with a piezoelectric transducer (reciprocal) by connecting the electrical terminals. In one direction, a mechanical input applied to the moving-coil transducer would produce a mechanical output from the piezoelectric transducer. If the moving-coil block is transformed to reciprocal by making force the flow and velocity the potential, then, at the electrical port, voltage will be potential and current will be flow. This is compatible with the electrical quantities of the piezoelectric block and the composite system will be reciprocal.

If, on the other hand, the two blocks are connected in parallel, we have a problem. Now both the mechanical ports are connected and both the electrical ports are connected. This could be done if both transducers were physically attached to the same structure or surface. The surface motion would drive both transducers. If force were chosen as potential to ensure

compatibility at the mechanical ports, the electrical ports would be incompatible: voltage would be potential for the piezoelectric transducer and flow for the moving-coil transducer. If voltage were chosen as potential for both transducer blocks, then the mechanical ports would be incompatible. The resulting combination could not be described by a reciprocal equation set. The resulting composite transducer would not be reciprocal.

The parallel connection is contrived to illustrate the point. Such devices are rarely encountered in practice but it's well to be aware that not every combination of reciprocal and anti-reciprocal blocks can be made reciprocal.

[Transparency III-4]

I'd like to dispense with a pair of myths about reciprocity. I was giving a talk a few years ago and someone in the audience from the calibration laboratory at this facility said they do not use reciprocity, because reciprocity works only if the transducer is lossless. His contention is not true. It's easy to come up with a simple two-port block that is lossy and reciprocal. This simple resistor network is one such block.

The block is linear and reversible. If you write the equations for the port voltages in terms of the port currents, you discover that B is identically equal C . The system is reciprocal but it certainly is not lossless; it is made up of nothing but lossy parts.

DR. GARRETT: I noticed that you omitted something twice when you spoke about the qualifications for a reciprocal system, so I want to know if it is intentional or inadvertent. You said that they must be reversible and linear but you never said that they must be passive.

DR. GABRIELSON: Good point. That was an inadvertent omission. A linear system with active elements (a transistor, for example) could be designed so that signals pass in both directions but it would, more than likely, not be reciprocal.

If you were analyzing a system with active components, though, as soon as you wrote the equations, you would find out that it wasn't reciprocal. The real crux is in the symmetry of the port equations. If the off-diagonal terms have equal coefficients, then the system is reciprocal: that is the definition of reciprocity.

DR. KEOLIAN: If it is reciprocal, you cannot say if it is passive or not. You can make an active system mimic a passive system and you could construct, in principle, an active reciprocal system and you could tell no difference between that and the passive system it mimicked. If you measure a system to be reciprocal, you do not know that it is passive.

DR. GABRIELSON: My feeling is that you are right. It ought to be possible to mimic a passive network by an active system so that its two-port equations are in the proper form and the off-diagonal terms are equal. If the off-diagonal terms are equal, then it is reciprocal.

DR. KEOLIAN: By definition.

DR. GABRIELSON: Yes. I am glad you brought it up. It is an interesting point.

We've seen that reciprocity does not depend on lack of dissipation. The other thing to keep in mind is that reciprocity doesn't guarantee that ALL quantities have equal ratios. If you put 10 V in one port and see 1 V on the other port, it does not mean that if you apply 1 V to that other port, that you would see 10 V on the first port. The ratio guaranteed by reciprocity is very specific. It involves the two quantities that I showed you on the previous slide.

DR. HOFER: I notice here you show only real impedances. What if there are complex impedances? Does that change anything at all?

DR. GABRIELSON: No. A passive network of resistors, capacitors, and inductors is still reciprocal as long as there is no nonlinear behavior. A ferrite-core inductor can introduce considerable nonlinearity so, in that case, all bets are off. In fact, most practical situations involve complex impedances.

[Transparency III-5]

Let's define some terms. Let's call the receiving response "alpha" and define it to be the voltage produced by some velocity input to this transducer. We will call the transmitting response "beta" and define it to be the force produced by the transducer for some input current.

Why is reciprocity useful? Reciprocity is useful because it provides a tool whereby you can replace difficult-to-measure quantities by relatively easy-to-measure quantities. In many cases, it is difficult to measure absolute force or absolute velocity. Voltage ratios are straightforward to measure. Electrical resistance, mass, and length are not bad to measure. What you try to do is contrive a situation in which you express the difficult-to-measure quantities in terms of quantities that can be measured accurately.

One of the often-cited papers in reciprocity has a title something like, "Absolute calibration of transducers without a primary standard." Technically, that title is okay, but it is misleading. You always need an absolute reference of some kind when you are doing an absolute calibration. What the author meant was that a microphone could be calibrated absolutely without having an absolute standard reference MICROPHONE. However, several physical quantities must be known absolutely to perform the calibration.

[Transparency III-6]

The classic configuration requires three transducers. There are cases where three are not required, but this is the simplest configuration to analyze. In the three-transducer configuration, one transducer is a source or transmitter, one is a receiver, and one is a reciprocal transducer. Of course, all three can be reciprocal, but only one is required to be.

There are three basic measurements that can be made. You can drive the source and receive with the receive-only transducer. You can drive the source and receive with the reciprocal transducer. And, you can drive the reciprocal transducer and receive with the receive-only transducer. In each case, you are going to be driving through some kind of transfer impedance. The ratio of the force produced at the receiver side of the apparatus to the velocity produced at the source side, for example. Or the pressure produced at the receiver to the volume velocity produced at the source.

The transfer impedance is really the key to reciprocity. The reciprocal of the transfer impedance is called the "reciprocity constant" in much of the reciprocity literature. This is unfortunate because it gives you the impression that there is something arcane involved. No. It is a transfer impedance, a very real, physical quantity that you can understand. If you can design the measurement apparatus so that the transfer impedance can be determined accurately, then you have a good reciprocity apparatus.

[Transparency III-7]

For simplicity, I've defined a ratio for each measurement and denoted that ratio by the source and receiver designations in square brackets. The ratio denoted [SR] means the ratio of open-circuit receiver (R) voltage to the input current to the source (S). The ratio [TR] means the ratio of open-circuit receiver voltage to the input current to the reciprocal transducer (T).

What are the unknowns? The unknowns are the source transmitting response (β -S), the receiver receiving response (α -R), and the reciprocal transducer receiving and transmitting responses (α -T and β -T). If you want to find, for example, the receiving responsivity, you multiply [SR] by [TR] and divide the result by [ST]. This gives an expression in terms of the desired quantity (α -R), the known transfer impedance, and also the ratio of β -T to α -T.

There were three measurements and four unknowns, so we need another relationship to finish the solution. If none of the transducers were reciprocal, we'd be stuck; however, the reciprocity relationship as it applies to the reciprocal transducer gives the fourth equation.

If you analyze an artificial case -- not an experiment that you need to perform but a paper analysis -- with the reciprocal transducer, you find out that α -T is identically equal to β -T. That equality is a result of my choice of definitions for α and β . If you choose those response definitions to be different forms, you will still derive the relationship between α -T and β -T but they may not be equal.

Since, with my definitions, the ratio of β -T to α -T is one, the measurements and the transfer impedance are all that are required to find the receiving response, α -R. In a similar way, you can find β for the source or α and β for the reciprocal transducer. You not only calibrate the receiving transducer; you calibrate all of them.

[Transparency III-8]

Here, I've listed some of the common variations of the "standard" three-transducer procedure.

[Transparency III-9]

Now, I'd like to go through some examples of reciprocity setups and show you what the transfer impedance is, since that is really the critical feature. If you read books on calibrating hydrophones this is what you usually see for the reciprocity arrangement. There is a source, a receiver, and a reciprocal transducer. The diagram shows only one of the three measurements. Here, the source and the receiver are positioned some distance away from each other and also far from any walls or reflecting boundaries. Some arrangements fix all three transducers in the water at the same time so that the ranges between different pairs are unequal but the essential feature here is that the transfer impedance relates to the propagation of sound from the source to receiver through some distance.

The transfer impedance is the pressure at location 2, the receiver location, produced by a volume velocity at location 1, the source location. This relationship comes directly from the expression I used earlier in the derivation of the mechanical impedance of the pulsating spherical source. The transfer impedance is the water density times the frequency divided by twice the source-to-receiver range. There is also a complex factor (with magnitude of unity) in front that is usually omitted from the literature because these calibrations are normally done to find just the magnitude of the receiving or transmitting response. However, the general reciprocity expression gives the phase of the response also and, if you are interested in the phase response, you have to leave that complex factor in the transfer impedance. Once you know the transfer impedance, you can make the three measurements and solve for whatever response you require.

If you need high accuracy, you wouldn't make the measurements for only one range. If you repeat the measurement set a number of times for different ranges, you can eliminate the

uncertainty associated with determination of the range. If the source is physically large, how do you determine source-to-receiver distance? Do you measure from the front face of the source? From the middle? From some other location?

You do not need to put up with that ambiguity. If you make the measurements at several different ranges, you can fit the response functions with a straight line and determine the response in terms of the increment in distance between measurement sets rather than in terms of absolute source-to-receiver distance. Increments are much easier to measure. As a byproduct, you are also able to determine just where the "acoustic center" of the source transducer is. Furthermore, the least-squares fit can reduce the influence of fluctuations introduced by boundary reflections in those cases when you can't move the transducers far from boundaries.

MR. WAKELAND: Thinking about what you were saying about absolute standards, even when you do this for several ranges, you still have to have an absolute standard distance to be able to determine the change in range between measurements, right?

DR. GABRIELSON: Yes. You need to be able to measure the distance increments absolutely. You need to measure frequency absolutely. You also need to know the density. You need to know the value of the constant, 2. (Laughter)

[Transparency III-10]

The next example is calibration of microphones in a pressure coupler, a very common way to calibrate microphones absolutely. You put the transducers into ports in the coupler two at a time (some couplers can accommodate all three) and make the required electrical ratio measurements. The transfer impedance is the pressure produced at location 2 for a volume velocity at location 1. If this system is much smaller than a wavelength, you can treat the coupler volume as a lumped compliance and easily calculate the transfer impedance. If the coupler volume starts to approach a wavelength, you can still calculate the transfer impedance but it's more difficult.

What do you need to know? You need to know the static pressure and the ratio of specific heats (γ) of the gas inside. You also need to know the coupler volume. A smart way to design the apparatus is to make the coupler tube with a central section or several central sections. By making the measurements with several volumes and least-squares fitting the results, you can reduce the calibration in terms of the volume increment of the central sections. This avoids having to determine the effective compliance from any of the transducers. (If one of the transducers is not completely rigid, then the acoustic compliance of the cavity is determined from the combination of gas-volume compliance and transducer compliance.)

Another useful trick is to change the gas in the coupler. Often hydrogen is used. Because the sound speed in hydrogen is so high, the coupler can be used to higher frequencies before wavelength effects become important. You have to be aware, though, that the viscosity of hydrogen is different than that of air so you will be changing the damping in the microphone if the microphone back vent opens into the coupler.

Sometimes you can set up a traveling wave tube. It is often hard to make a good traveling wave but, if you can, then the transfer impedance is simple: gas density times gas sound speed divided by the tube cross-sectional area.

MR. WAKELAND: I am under the impression that in order to do a reciprocity calibration you have to set up a geometries that you know and understand. You could not just put three transducers, some of which are reciprocal, inside a car and then do a reciprocity calibration.

DR. GABRIELSON: You could but you would need to work pretty hard to find out the appropriate transfer impedance. Simple geometries lead to simple transfer impedances, but, if there was some pressing reason to do the calibration in a car, you could, in principle, determine the transfer impedance.

DR. GARRETT: If you had a calibrated transducer, you could measure the transfer impedance in the car.

MR. WAKELAND: Yes, but if I had calibrated the transducer, I could just go back and do the comparison calibration.

DR. GARRETT: Let's say you want to calibrate a set of transducers and the conditions will be changing. You take very expensive calibrated transducers and you establish a transfer function. Now you remove the expensive transducers and put cheaper transducers in their place. Now that you know the transfer function, you can let the cheaper transducers calibrate each other.

DR. GABRIELSON: Right. Another thing I did not mention is that the transfer impedance may be different between different pairs of transducers. This is a straightforward change. Go back to *Transparency III-7* and put different transfer impedances in each of the three measurements, write the ratio expressions ([SR], [ST], and [TR]), then write the expression for $[SR][TR]/[ST]$. Solve for α_R . The expression is slightly more complicated but nothing fundamental has changed.

[*Transparency III-11*]

Calibration inside a rigid-walled resonator is a nice problem to analyze. My favorite reciprocity paper is the one in your reference material by Isadore Rudnick. Don't be misled by the title, "Unconventional reciprocity calibration." The paper deals with the fundamental roots of reciprocity calibration. Prof. Rudnick talks about the rigid-walled resonator in this paper. He leaves several steps in the derivation to the industrious reader and I have filled in the missing steps in this slide. In this case, you get a transfer impedance that is a function of the density and sound speed of the gas inside the resonator, the Q of the resonator, and the cross-sectional area, A, of the resonator. Presumably, if you had calibrated transducers and you knew the density, sound speed, and resonator Q, you could determine " π " experimentally. (Laughter)

[*Transparency III-12*]

We've been talking about pressure transducers -- microphones and hydrophones -- so far in these examples. Accelerometers are straightforward as well. If you put accelerometers on a pendulum and at least one of them is reciprocal, you can use one to drive the pendulum. Supplying a current into the drive transducer produces a force on the pendulum mass. The pendulum moves so the receiving transducer moves. The transfer impedance is " j " times omega times the mass.

The transfer impedance is simple. It is just the mechanical impedance of a free mass. For better accuracy, there are a number of corrections that can be made. I wrote a paper on this system that discusses the more important of these corrections. For example, the dynamics of the transducers themselves affect the impedance.

Also, if the frequency is high enough, the pendulum mass looks more like a section of a waveguide than a simple mass. Even without these corrections, though, reasonable results are possible with the simple mass-like transfer impedance.

For sensors that respond to rotation or changes in rotational rate like gyroscopes, the linear pendulum concept can be translated into a torsional pendulum. The transfer impedance is the impedance associated with the moment of inertia.

DR. FORTUNKO: Some of the pendulums that we have, have hysteresis in them. Is that reciprocal?

DR. GABRIELSON: That would be a serious problem. The pendulum I've shown is suspended by strings and the motion is very clean. If the suspension has hysteresis, you would not have a reciprocal system.

DR. MIGLIORI: It is not linear.

DR. GABRIELSON: Right. The suspension is critical. Simple, but critical.

[Transparency III-13]

This may be the simplest version of the pendulum calibrator: three transducers clamped in a fixture suspended by four strings. Notice that the suspension is designed to keep the pendulum motion in a plane containing the accelerometer sense axes (the direction of the arrow labeled "v"). It is also arranged so that the beam containing the accelerometers does not rotate. If the beam rotated, then its angular momentum would need to be considered in addition to the linear momentum. I use a better fixture -- a solid cylinder with screw mounts for an accelerometer in each end. This works better at higher frequencies and the waveguide correction to the mass is much easier to calculate but the arrangement shown is simple and sufficient for demonstration. This is a nice teaching lab apparatus.

[Transparency III-14]

You do not have to use reciprocity as an absolute technique, either. Here I've shown an example of doing a comparison calibration by reciprocity. If you drive a reciprocal transducer to force the pendulum and measure the motion of the pendulum with a calibrated accelerometer, you can deduce the receiving sensitivity of the reciprocal transducer by reciprocity. Sometimes it is easier to do that than mount both transducers on a shaker or to do a full three-transducer reciprocity calibration. Geophones excellent reversible transducers for these sorts of experiments, by the way.

[Transparency III-15]

I want to talk now about another technique for absolute calibration. When I started developing the pendulum techniques for reciprocity calibrations, Professor Garrett talked me into doing a calibration by optical interferometry on the same apparatus. That would give me two different absolute measures of the same quantity.

This is a good general philosophy in measurements. Reciprocity is reliable, but, it's very comforting to have another means of calibration that is independent of the parameters used in the reciprocity calibrations.

The example that I've shown is one of many setups. I actually have a free-air-path interferometer in my lab at Penn State, but when I did the original work, I used a fiber-optic interferometer system. A laser feeds a coupler-splitter; half of it goes into one of the output

fibers and is "stopped" by a small amount of cement or oil on the end of the fiber so that the light is not reflected. Shattering the end of the fiber also works.

The other half goes into the "sense" arm. The light is partially reflected from the nicely cleaved end of the fiber and part of it leaves the fiber, bounces off of the surface that you are measuring, and re-enters the fiber. These two paths interfere and the interference is detected by the photodetector after the light passes back through the coupler.

[Transparency III-16]

The component that reflects from the moving mirror has a phase that depends on the motion of the mirror through the "two k -zero d " term. The photodetector produces an output proportional to the incident optical intensity not the incident electric field. The intensity depends on the square of the electric field. Because the optical field oscillation is at a frequency much higher than the detector electronics can track, the resulting output is proportional to the time average of the electric field squared. This is just like a square-law detection process. The two interfering components add coherently but the detector produces the time average of the intensity. Consequently, the photodetector produces a signal proportional to the cos of two times k -zero times d .

If the mirror is moving sinusoidally, then d would be some reference displacement, d -zero, plus some amplitude, d -one, times the sine of the vibrational frequency of the mirror.

[Transparency III-17]

If you have carefully set up the optical path lengths so that the equilibrium displacement is at the maximum rate of change point of the $\cos(2kd)$ function, then a small sinusoidal displacement of the mirror maps into a sinusoidal output.

With a spectrum analyzer you would see a line at the frequency driving the system. If the path lengths drift a little bit, then the response changes. In fact, if the equilibrium point moves to a peak of the $\cos(2kd)$ function, the nonlinearity is so strong that the primary output is at twice the frequency of the input.

[Transparency III-18]

Fortunately, it is not necessary to limit the excursions to small amplitudes about the "proper" equilibrium point. If the mirror moves far enough, then the output looks kind of crazy, but it is well-described mathematically.

If you look at it on a spectrum analyzer, you would see a line at the frequency of excitation, a line at two times that frequency, a line at three times that frequency, and so on.

The amplitude ratios of all of the lines are given by Bessel functions, J_1 for the line at f_1 , J_2 for the line at f_2 , etc.

[Transparency III-19]

In fact, if you record the output of the spectrum analyzer at one of those lines as you adjust the drive amplitude, you map out a Bessel function. As the drive amplitude increases, the f_1 line first increases, then peaks, then decreases to zero, then increases again just as the Bessel function, J_1 does for increasing argument.

Why is this useful? If you work out the math, the location of the nulls does not change if the equilibrium point changes. You don't have to tune the system to the maximum sensitivity point and try to keep it there. As long as you are creating large enough displacements, you can find the drive level that causes one or another of the lines to null out. Of course, for this

procedure to work, you need to have enough displacement to get to at least the first null. Especially at high frequency, that can be difficult. But with transducers capable of producing displacements of the order of an optical wavelength or more, the technique is very useful.

[Transparency III-20]

The form of the function that describes the photodetector output at frequency 1, for example, is the product of a sine function and J_1 . If you are looking for zeros, the behavior of the sine function with its equilibrium displacement argument is immaterial unless you happen to be at or very near the point where it is zero.

The pendulum reciprocity calibrator is always moving several optical wavelengths because of room vibrations so zeros of the sine function are passed through very quickly. In a more stable, fixed system you might accidentally wind up near a zero of the sine function and that would degrade the signal-to-noise ratio.

Having dispensed with the problem of zero response caused by a poor choice in equilibrium path difference, we can focus on the zeros of the Bessel function. Notice that these zeros depend only on the displacement of the surface and the wavelength of the laser. If you find a zero crossing, the error is primarily determined by your ability to localize the null. You would often know the wavelength of the laser to much greater precision than you would be able to exploit.

DR. LAUTERBORN: Won't you get something else if you use an excitation that is not sinusoidal?

DR. GABRIELSON: Yes. Now you don't know how to interpret nulls even if they do appear in parts of the spectrum. For calibration, you would use sinusoidal excitation. It's another limit of the optical technique. Reciprocity, on the other hand, can be done with nonsinusoidal excitations (random noise, for example) as long as the response remains linear.

The pendulum calibrator is an interesting case. Normally, you try to arrange an interferometry setup to be very stable. The pendulum, however, has a natural resonance at about 0.6 Hz with a Q of over 500. Building vibrations couple in to the pendulum through its suspension and those vibrations near 0.6 Hz are Q -multiplied so that the pendulum swings with an amplitude of many wavelengths at that frequency.

In a linear system, like the reciprocity measurement, this oscillation at 0.6 Hz is only a problem if you try to push the calibration frequency down below 1 Hz. In contrast, the interferometer produces a nonlinear modulation. For large (in terms of wavelength) swings at 0.6 Hz, the harmonics produced can extend far up in frequency. I have a great deal of difficulty performing the optical calibration below about 50 Hz because of these harmonics. In fact, the only way I can reach down to 50 Hz with good repeatability is to put an eddy-current damper on the pendulum to drop the Q , which reduces the level of oscillation at 0.6 Hz.

DR. FORTUNKO: This technique works because you drive the pendulum many wavelengths, right? If you were to use an ultrasonic wave at megahertz frequencies, you might not get the displacements necessary to drive it that far.

DR. GABRIELSON: Right. I've just discussed the low-frequency limit to the technique as applied to the pendulum. The high frequency end is a problem because it becomes harder and harder to push the pendulum far enough to get to the first null. With the transducers I was using, I was having trouble reaching the first null beyond about 400 Hz, so I had a pretty narrow

window in which I could do the Bessel-null calibration. A much smaller pendulum mass would permit higher frequency excitation to the first null.

This technique tends to be more restrictive than reciprocity, but, if it does work, it is a nice technique. The zeros of the Bessel function are well known. You can easily find a laser with a well-established wavelength. I was using a helium-neon laser with more decimal places in the wavelength than I would ever use.

In my case, it was particularly useful because the optical results do not depend on the same things that the reciprocity results do. The optical method depended on knowing the wavelength of the laser light and being able to measure absolutely the voltage output of the receiving transducer. The reciprocity calibration depended on the resistance value used to convert the drive current into a voltage, the mass of the pendulum, and the dynamical corrections to those masses.

DR. GARRETT: The value was because the bandwidth of the two calibrations overlapped.

DR. GABRIELSON: Yes.

DR. LAUTERBORN: You could also shift the upper frequency limit by using a laser of blue light, for instance.

DR. GABRIELSON: Yes, if you use shorter wavelength light, a smaller displacement reaches the first null.

[Transparency III-21]

I would like to close by reminding you that a sensing system is usually more than just a single component. When you are analyzing a sensing system, you have to look at all of the parts. Some people in the transducers business think that a transducer is always just a piece of piezoceramic with two wires connected to it, but there are many kinds of transducers that are more complex than that (and some that are simpler).

A capacitive microphone, for example, usually has a DC bias on it and the bias can introduce noise or affect the calibration. Less commonly, a capacitive microphone can have an AC bias. There are high-precision capacitive accelerometers that use high-frequency AC bias.

Some systems have feedback: the proof mass is forced to maintain a constant position with respect to the case instead of being allowed to move. The signal applied to the balancing actuator becomes the sensor output. More acceleration of the case requires a greater force to keep the proof mass stationary with respect to the case; a greater force requires more signal to the balancing actuator; consequently, the output signal is higher. This feedback process contributes to the overall system noise.

Is the feedback control done as an analog process; is it done as a quasi-digital system (a sigma-delta converter, for instance); is the feedback sensing done optically? All of those approaches have impact on the system noise and on the system calibration.

You should know what the ambient or external level of noise is for your application environment. It does not make much of sense to design the transducer to be twenty times quieter than the quietest background you would ever experience.

Is the sensor responding to the desired signal and not responding to other excitations? You might have a sensor that is exceptionally sensitive to acoustics but it is also exceptionally sensitive to temperature. Often, that's not good.

There are at least two philosophical positions to avoid. Avoid believing that, because you've used a particular technique many times and you like it, there is only one best technique to use. Be flexible in your thinking. Study other approaches. Study sensors in fields other than your own.

Try also to avoid the tendency to "solve" signal detection problems by assuming that sufficiently powerful digital signal processing will pull the desired signal out of the background noise. There is a great deal of power in DSP these days, but intelligent design of the sensor and the measurement apparatus can pay real dividends, even if you are eventually going to use DSP.

DR. FORTUNKO: We just presented a paper where we showed that most of the ultrasonic digitizers are 8 bits. We showed that that can really get you into trouble, because of digitizer noise.

DR. GABRIELSON: Good example. Thank you very much. My time is up.

[Sensor Electronics Supplement]

THERMOACOUSTICS MADE SIMPLE

Steven L. Garrett
Graduate Program in Acoustics
Pennsylvania State University

With Demonstrations by

Thomas J. Hofler
Department of Physics
Naval Postgraduate School

[Transparency 1]

DR. GARRETT: Thank you. This could be a little tricky, because I am going to turn the lecture over to Tom Hofler in a couple of minutes for some excellent thermoacoustic demonstrations. Before I start, I would like to thank the organizing committee, Hank, Logan, Anthony, Libby, and Logan, again, for picking up the tab. It is a real pleasure to be asked to come back here and lecture. Of course, I am always asked to lecture on something different; I screw up one thing and then the next time, I lecture on something else.

Before we get to the course material, I want to make some general about summer school. On the first day you were told a little of the history, but the history goes back a good deal farther.

[Transparency 2]

I was very fortunate in 1974, to be able to attend the Physical Acoustics Summer School held as part of the Enrico Fermi Summer School. The Italian Physical Society has a villa, and this is a real *villa* -- leather tooled wallpaper, marble statues of guys holding grapes along the verandah overlooking Lake Como, a beautiful lake on the Italian-Swiss border.

The Italian Physical Society would get this villa for 2 months and they would sponsor four 2-week summer schools. I was fortunate enough to attend the one in 1974. Professor Quate talked about the one that was held in 1984, which, unfortunately, I was not able to attend. There was an earlier one in 1964, and there has not been one since the 1984 summer school.

The reason that I put up this school photograph is that there is a tradition in physical acoustics that the old guys get together in some lovely location and bring their grad students and do what we are doing here. When I was a grad student back up there, I never realized that this was a tradition that was going to get passed on and that I would have a responsibility for its continuation. (Sort of a science education "chain letter.")

For those of you who are as dense as I was, I put this up, because you will have the responsibility 20 and 25 years down the line to make your contribution. At that time, the "old guys," the front line in the photo, were very distinguished acousticians, so you guys have a little bit of a disadvantage here in Asilomar. (Laughter)

This guy in the lower photo at the far left is R. Bruce Lindsey. He was the editor of the *Journal of the Acoustical Society of America* since almost year zero. Next to him is Warren P. Mason. He was Mr. Acoustics, Mr. Ultrasonics, at Bell Laboratories. Daniel Sette is a professor an ultrasonics researcher in Italy, a very famous guy in Europe. He was the one who would get

us into the queue for Varenna (he had the "in" with the Italian Physical Society). This is my esteemed adviser, Professor Rudnick.

I think this is Claus Dransfield who is a pioneer in very high frequency ultrasonics for the study of phonons in solids. Richard K. Cook is probably best known for his invention of the reciprocity calibration technique and for some amazing early studies of infrasonic sound propagation over continental distances. Edward F. Carome, at that time was studying piezoelectric polymers and then became involved in the development of fiberoptic sensors shortly thereafter.

Anyway, I wanted to call your attention to the fact that this type of summer school is a tradition in physical acoustics. I think it is a very good tradition. I learned an awful lot by watching these old guys work.

A few more people worthy of note are shown in the upper photo. For the Penn Staters in the room, this is Ralph Goodman. I showed him this photo before I left for California and he took pleasure in pointing out that he was about my age at the time that picture was taken. Taylor Wang was another of Izzy Rudnick's graduate students. He flew on the *Space Shuttle* (STS-24, the one before the Challenger accident) to do some experiments in acoustic levitation.. Sid Yip, a professor at MIT. John DeKlerk, a South African, very famous in surface acoustic-wave devices.

This was really a great selection of outstanding acousticians who would sit around under Cinzano umbrellas and criticize each other's work the way we do.

I am going to start the talk by having Tom Hofler show us some demonstrations that he has built at the Naval Postgraduate School. Before I do that, I would like to do what Anthony did at the start of his talk. I would like everybody who is a "thermoacoustician" to please rise. [Many participants rose.]

I am now going to ask those of you standing to please leave. (Laughter)

I will do just like Anthony; there are people here who are competent in areas, many areas, more competent than I am and I will freely pass some questions and explanations to them as this lecture progresses. The first pass will be to Professor Thomas J. Hofler. A lot of you stood up a moment ago. In 1984, when Prof. Quate was talking about John Wheatley giving his thermoacoustic lecture, there was exactly one graduate student in thermoacoustics: Tom Hofler.

He was the first one to put together an actual device where one could make careful measurements, both thermally and acoustically. I was very fortunate that when he received his Ph.D. he agreed to come and postdoc with me and, it is true, I learned a tremendous amount -- not enough, but a tremendous amount -- working with Tom and our graduate students building a device that we worked on together for many years.

I would like to ask Tom to take over. I think when you see what he is going to show you, it might get you through the drier parts of my lecture.

DR. HOFLER: Thank you, Steve. Steve was gracious enough to let me bring a couple of demos out. Steve also has nice demos, but they are at Penn State and you have the troubles of getting them on airplanes, and so on, so it was easier for me to bring mine over from Monterey.

As Steve mentioned, we are doing the demos first, which is usually the reverse order. Usually when I do these lectures, I talk about theory and measurements first and then do the demos at the end. I think that is the way Steve usually does it, too, but it is a little livelier this way.

I brought two demos, one that is a thermoacoustic refrigerator and one that is a thermoacoustic prime mover, as it is sometimes called, although one of the fathers of thermoacoustics, Nicholas Rott, said he hated the term "prime mover." I kind of like to call it the thermoacoustic motor. If you think in terms of heat engines, this is really a motor.

I have 2 demos and I am going to actually reverse the usual order here. I am going to show you the refrigerator first.

[Transparency H1]

This is the most mechanically complicated of the two, although from Steve's lecture on the theory, personally I find the theory a little simpler for the refrigeration. The prime mover, which I will show you next, is mechanically much simpler, but I think the theory is actually a little more complex.

By the way, this is a section engineering view of this thing right here. Let me just show you the parts here, and I will start from the bottom. There is a fan that drives airflow through this to cool the heat sources, which are the electrodynamic driver. This sucks airflow through an air duct and not only cools the heat dissipation of the electrodynamic driver but also the heat rejected by the refrigerator that is at the hot heat exchanger -- here.

Moving up, we have part of the pressure vessel for the driver. You can see that a lot of this bulk you are looking at is actually the air duct, and the fan and the driver is fairly small. This is the magnet and the pole pieces on the other end, a moving coil right here that is reduced to a smaller diameter, a vibrating piston -- we drive it around 650 Hz. The vibrating piston drives this whole quarter-wavelength resonator here.

The driver piston is at a pressure anti-node in the standing wave, so the velocities are quite small here. This is effectively the quarter wavelength, roughly speaking. This is the velocity anti-node, about right up here.

Another way to look at this standing wave in here is to think of it as what Mike Moldover likes to call a double Helmholtz resonator, where you can think of this as being one gas-compliant volume, this being another gas-compliant volume up here, and then the small diameter tube is the gas mass, and this mass then bounces against two gas springs.

Of course, we have computer models for this and the real calculation is considerably more detailed; I will not go into that.

All the components are important, but in terms of the thermal components here, and the thermoacoustic components that are critical, that is really right here in this region. As you move upward from the piston, the velocity increases some. You have what we call the hot heat exchanger (it is actually slightly above room temperature -- you could call it an ambient heat exchanger) and above that is a stack.

The word "stack" is just a short way of referring to what was originally a stack of parallel plates, so stack is short for stack of parallel plates. Of course, nowadays it is not actually parallel plates; it is a spiral roll. This is actually very similar to what is inside the demo 'fringe.

[Transparency H2]

This is not quite a full diameter stack with leftover material, so you can get an idea of what we are talking about. That is Mylar plastic film or polyester and we glue fishing line onto the surface of the plastic film on a long strip and then roll it all up.

Let me go ahead and turn it on. I have got an LED thermometer here reading in degrees Centigrade.

[The refrigerator was demonstrated.]

Unfortunately, this is a little bit noisy and you also hear the whir of the fan, but it is starting out at about 22EC. You can see already the temperature is dropping here, and I will talk about this for a couple more minutes while it runs and then I will turn it off.

As you can see, the temperature is dropping. The thermometer is reading the temperature of the thermocouple attached to this part right here, the cold heat exchanger; this whole copper part.

DR. MIGLIORI: What is the horn for?

DR. HOFER: I talked to a Stirling engine guy once and he said, "Oh, that's called a diffuser." The idea is you have fairly high velocities and Mach numbers here and the obvious thing to do is simply cut the tube off at some point. What you are doing is you are ejecting the sound wave into a large open volume at a fairly high velocity if you do that.

The idea of the little horn and the flare at the end is to try to slow the acoustic velocities down a little bit and provide a little flare so as to avoid lots of flow separation and turbulence or vortex shedding at the end of the tube, which would be a nonlinear dissipation mechanism.

PARTICIPANT: Is the shedding of vorticity sensitive to the angle of the horn or the flare?

DR. HOFER: That is something we have not really investigated a lot. I really got the idea from looking at what engineers typically do with plumbing for steady flows; higher velocity water flows and gas flows. If you look at ordinary plumbing fittings, a lot of the better ones will have flares in them specifically to avoid these kinds of vortex shedding and so on, which would show up as an extra flow impedance.

I have not really investigated the details of the acoustics but I assume if it is important for D.C. flow, it has got to be important. There are papers, old JASA papers. Sound through an orifice, there are 3 or 4 regions of phase space where you go from laminar flow to vortex shedding to turbulence to jet formation; so these things have been mapped out.

In fact, Mike (Moldover) and his coworker have done work on double-Helmholtz resonators, where they chamfered the ends of the pipe. They find, even at low amplitudes, a significant change in the dissipation.

PARTICIPANT: I just wondered whether it is important to completely stop separation happening from the edge -- in other words, vortex shedding from the edge of the pipe. I thought maybe there might be a critical angle or something, depending on the Reynolds number.

DR. HOFER: There may be some such complicated dependencies. In some rare cases we can drive it hard enough where we see what looks like the resonator dissipation becoming nonlinear or excessive. It seems to be that we are operating below that amplitude most of the time.

MR. WAKELAND: The big-brained people who were at the TRITON design meeting, including Swift, I believe, said that they had gotten the impression that 7° was the most you could flare out without getting - - -

DR. HOFER: That was the number that I had read about way back in the early 1980s when I made the first one of these kinds of resonators, and it was just standard engineering practice for steady flow in pipes.

DR. GARRETT: May I comment? The way I like to answer that question is to point out that locomotives were running at 100 mph in the U.K. when Joule wrote down the First Law of Thermodynamics. A lot of this stuff is being built by intuition. We will probably try to study this detail, because it is important -- but in the early days of a new technology, you build it and keep your fingers crossed. We are at that stage of the technology development where those answers are probably known in some other field or not known and you don't have the time to find out which is the case.

DR. HOFER: You have the stack here, your cold heat exchanger, which I have mentioned. This is helium gas and the helium is at 90 psi absolute in there.

Let me pause right now and bring Craig Hodgson up. Craig is going to give me a hand on the next demo and it needs to be pre-cooled. Craig has actually built one of these himself, one of these Hofler tubes, which I will talk about in a minute. Since he has built one, I thought I would have him lend me a hand. He is going to pre-cool it now.

I want to mention that we had three masters thesis students working on this demo 'fridge. One built the modified electrodynamic driver for purpose of our cryocooler experiment. We built an extra one and had it lying around, so we decided to build this demo refrigerator as something that we could do with the extra driver.

We were constrained by the limits of the driver, which had an upper frequency limit of about 650 Hz. I then had 2 more thesis students who built the rest of the apparatus. Our objectives were to try to make it as compact as we could, given the frequency running in helium, so we could fit it into a case and sacrificed some efficiency for that reason.

The other motivation is that in a lecture like this I can show you a thermometer and you can watch the digits change, but that is not very exciting. After the talk I am going to turn this on again and you can come up here and since the cold part here is exposed. You can touch the cold part and feel the cold, which is maybe a little more impressive.

Then we thought, well, gee, if we work at it, maybe we can get this thing to frost up. In fact, depending on how warm it is in the room and how much breeze is blowing by, since this is exposed here, usually it will collect frost. There is a safety warning for this and that is, do not stick your tongue on the cold heat exchanger.

[Transparency H3]

Let me quickly show you a little bit of data. These are raw data that we took in several different places on the refrigerator. You can see that we get down below 0°C. A couple of other interesting points are that this curve here is actually on the top of the resonator and that actually does heat up, at least initially. While the cold part is cooling down, the top of the resonator does heat up from acoustic dissipation.

In addition, you notice there is a little spike in the cold data here and at about -3° to -4°C. That is where you are making ice. We actually tried to get the humidity out of this measurement, but have not done quite a good enough job and that has a really dramatic effect on your data.

[Transparency H4]

The next thing I am going to show you is based on essentially these raw data. These are reduced data, based on the raw data from the last slide. We fit these exponential curves, measured the weights, calculated the heat capacity of our components, and inferred heat power

from the exponential temperature change and the heat capacity of the parts. I would not call it a really high-precision measurement, but it is good enough for the demo fridge.

PARTICIPANT: Was that water vapor in the air?

DR. HOFER: Yes.

You can see that the ultimate temperature span is about 50°C. This is a coefficient-of-performance that Steve will talk about, hopefully -- it is not an efficiency, but you can think of it like that -- versus the cooling power down here. This has a COP of about 1, if you are several degrees below 0°C.

I think Craig is ready to go here, whether I am or not, so hang on. Craig has this thermoacoustic prime mover ready to go. He is going to pull it out of the liquid nitrogen and you will hear a loud sound made thermoacoustically. I will squirt helium in the open end of the resonator and it should be louder. Then maybe we will put it in front of the projector to see the jet coming out of the end. [Demonstration of the prime mover]

It always impresses me that something like flimsy tissue paper can kill the whole effect.

[Transparency H5]

Let me give you a quick picture of what is in the Hofer tube; it is really pretty simple mechanically. It is basically a pipe that is closed at one end here and open at this other end here, so the fundamental mode is a quarter-wavelength.

As usual, the important stuff is all in the middle here, focused around the stack. You have a hot heat exchanger that, again, is at ambient temperature, another spiral roll- plastic film stack, and a cold heat exchanger. There is a thin-wall stainless steel tubing right here, and these copper tubes and copper flanges here conduct the heat out of the cold heat exchanger fins, the flange, and the tube wall, and then into the liquid nitrogen at this end. Of course, your hand or the ambient temperature keeps this end warm. Again, the thin-wall stainless steel tubing is a pretty good thermal insulator.

It is really not very difficult, mechanically very simple, but as I said, the theoretical explanation is a little bit more involved.

With the refrigerator demo, it is obvious what a refrigerator is good for. People want to have refrigeration. This Hofer tube prime mover demo is -- you can say, well, gee, it is really cool, you make a loud sound, but what is it good for? Well, it is an acoustic motor.

[Transparency H6]

I want you to focus just on the figure here. You can combine an acoustic motor and, instead of using liquid nitrogen to suck heat out it, you use a high-temperature heat source to force heat into it. Essentially the upper half of this thing is a large Hofer tube, where the sphere is closing off what would have been the open end. Of course, the other end of the Hofer tube has an end plate or end cap.

Because this is now a half-wavelength design from a velocity anti-node to a pressure anti-node to a velocity anti-node, you do not need an end cap. In fact, the pressure anti-node is formed naturally in the middle, so the upper half is the thermoacoustic motor and the lower half is a thermoacoustic refrigerator.

[Transparency H7]

This is what it looks like in the laboratory. I wanted to show this partially because this is Jay Adeff, who is the technician-engineer working for us and has done most of the engineering and building all these things.

[Transparency H8]

Finally, we got it to work and produced about 100 W of cooling power. We plumbed the cold antifreeze into a Styrofoam cooler and here I am drinking a cold soda from this cooler that was chilled thermoacoustically with no moving parts -- heat-driven refrigeration.

That is why acoustic motors are nice, one of the reasons you can do heat-driven refrigeration with no moving parts and no loudspeakers.

DR. GARRETT: Thank you very much, Dr. Hofler.

DR. HODGSON: I would like to interject that building with Hofler tubes it was quite easy to put together and it makes an excellent demo. You really grab people's attention, even when they are walking down hallways. It is quite fun.

DR. GARRETT: I should have given you that warning in the beginning -- thanks, Craig -- this stuff is addictive. It does not leave needle marks, but is definitely addictive.

[Transparency 3]

What I hope to do this afternoon is start out with the motivation. Basically, thermoacoustics started as a curiosity about the interaction of sound and heat and a desire for understanding at a very basic level of what is going on. It would have stayed that way had over one hundred nations had not ratified the Montreal Protocols. I will make some statement about the impact of global environmental treaties and their effects on this research. Thermoacoustics is attractive both because of what Tom showed you and because it is a very simple device that works quite well -- in fact, we think it works best with inert gases and, therefore, it has no adverse environmental effects.

There is another good reason for doing thermoacoustics. Professor Quate, I think quite eloquently pointed out that you go into a university and you are faced with having to think of what to do to educate graduate students. I was very lucky in working with Professor Rudnick. At that time, superfluids were an interesting area of research that trained us in hydrodynamics, trained us in instrumentation and calibration, construction techniques, in cryogenics, and other technologies. It also had good acoustics in it and good engineering challenges.

Albert (Migliori) was educated as a graduate student through the superconductor research area and we have seen several people here who came up through cavitation research. I would say that, irrespective of whether or not you find thermoacoustic refrigerators in the basement of Sears 20 years from now, we will have trained a lot of grad students in a very challenging and a very broad problem, so I have a dual motivation for enjoying this research problem.

What I plan to do is cover some of the very basic fundamentals. I hope I do not insult too many of you -- remind me to ask when I move on to standing waves. I am going to talk about the things that show up in some of the other lectures you have heard.

We are coming to the end of Summer School. Tomorrow is the last full day. If physical acoustics is really "a field," there should be a lot of overlap. A lot of what I am talking about should be linked to Mark Hamilton's nonlinear lecture and to Jim Sabatier, who kept bringing in the viscous penetration depth/thermal penetration depth concept. Besides, I think that is probably worthwhile looking at these fundamentals.

One of the reasons that many of the people who worked in one area as grad students and are now in other areas are able to contribute in other areas is that we had a good handle on the fundamentals. Those fundamentals do not change; the applications change. I am going to stress this as much as Anthony did when he talked about just bubble dynamics. You really want to understand the part that is known before you go on to the unknown acousto-optic conversion mechanism: sonoluminescence.

I am going to apply these thermoacoustic fundamentals to a very simple standing wave. I will give you what I believe will be the most important, possibly only important, viewgraph in this talk. It is a Lagrangian (moving gas parcel) model. It is a good physical picture of what is doing the heat pumping and what is generating the sound, without worrying too much about the mathematics (the mathematics is going to be extraordinarily sloppy).

I can get away with this sloppiness, because there is an excellent detailed model, and I will show you some data from Reh-lin's thesis measurements, and Matt Poese's thesis measurements will show you that we really do understand some things about thermoacoustics at a very, very detailed level.

I hope I can just give you a basic intuitive feel for why there is a "critical temperature gradient," why there is heat and energy transport in a standing wavefield, and what effects sound waves in gases near solid surfaces. I will argue that I have been successful in this field because it is superficial; everything happens on surfaces. (Laughter)

I must say, that per-unit volume, Jim Sabatier is more superficial. He has to deal with more surface area per unit volume in his work, so I do not feel too bad about my own superficiality.

We will look at some of the limitations imposed by the first and second laws of thermodynamics. Dr. Hofler pointed out that COP, coefficient-of-performance, is an important metric and we will have to get that easily under our belts.

I will talk about some acoustic refrigeration systems that are actually practical, or we tried to make them practical. I will finish up in whatever time remains discussing some of the engineering issues, some of the components that dictate how well something is going to work and how one attempts an optimization, how one makes at least initial choices such as to where to place the stack in the standing wave.

Tom had his stack there. You do not know why. The prime mover had a stack somewhere else and it was a certain length, and I want to bring some of these things out. How do you design a loudspeaker? What are the issues that control piston size and efficiency? In addition, show you some of the directions that we are headed in our group.

CFCs, chlorofluorocarbons, are responsible for a major change in the way people live. It is not as big as the automobile, but it is mighty close. There are people who have said that the invention of refrigeration saved more lives than the introduction of medicine. More people used to die of food poisoning than died of polio and smallpox, *et cetera*, over time. There are certain parts of this country that would be uninhabitable if it were not for refrigeration.

By 1996, when they stopped manufacturing CFCs, \$40 billion worth of CFC-filled equipment, air-conditioners and refrigerators, were sold in the U.S. alone; every year, \$40 billion a year! It was an enormous industry.

It was postulated in 1972 that the chlorine from the CFCs was causing damage to the Earth's stratospheric ozone layer. The immediate response from the chemical industry is what I

call the tobacco defense: It "ain't" there. As you know, Rowland and Molina got the Nobel Prize in chemistry two years ago for showing that it was there.

The refrigeration and chemical industries then moved to what I call "the Detroit defense": If we lost this industry, we'd lose a lot of jobs and, besides, it can't be done any other way. Eventually what has happened is those alternatives chemicals have been developed that do not contain chlorine, which is what catalyzes the destruction of the ozone, or contain very little chlorine. These are called HCFCs, hydrogenated chlorofluorocarbons, and HFCs, hydro- -- whatever. Chemistry!

Never trust a skinny chef. Never take chemistry advice from a guy who did his thesis on waves in helium. Helium has no chemistry. (Laughter)

So give me a little slack on the chemistry, but the point is that the HFCs do not destroy the ozone, but they are over 3000 times more potent as a global warming gas than CO_2 . The Montreal Protocols were signed in 1986. The ban went into effect in 1996. The Kyoto Accords, which are in the process of being ratified (China signed about 2 weeks ago), propose to charge manufacturers for "carbon credits."

Eventually one would like to leapfrog this entire cycle of chemical dependency and go with something that is chemically inert. Thermoacoustics is certainly one of the leading technologies.

DR. MIGLIORI: Steve, whatever happened to sulfur dioxide and ammonia, which now look environmentally benign?

DR. GARRETT: If I get into this thing, we are going to lose a lot of theory. Ammonia has got problems because it is so caustic. The CFCs were developed as a substitute for ammonia. The Germans are tending to go with hydrocarbons, which are flammable and explosive. There are other alternatives. Thermoelectrics are an alternative. They are a bad alternative because they are so inefficient that you end up generating more CO_2 to make the electricity to run them.

Anyway, I just wanted to establish the fact that I have gone, after working for the Navy for 12 years, from being a "tool of the military industrial complex" to being an "award-winning environmentalist". Life is full of surprises.

[Transparency 4]

Shall I get to the science? Thermoacoustics is just what it says. It is "thermo" and "acoustics;" the interaction of heat with sound. The problem of the interaction of heat and sound is as old as the field of mechanics itself.

The first fundamental calculation of sound speed was done by Isaac Newton, published in *Principia*. Newton showed from theoretical arguments that the speed of sound in air was 979 feet per second. It turns out when they measured it, he was wrong. He hemmed and hawed, criticized the experiments and it took another hundred years until Laplace solved the sound-speed problem it correctly.

Let me just show you the ideal gas law (1): $pV=nRT$. I have written the number of moles as the mass of the gas, m , over the molecular weight of the gas, M . I could stop here and hand everything over to Mike Moldover, but I am not going to do it.

You have seen many times that the sound speed, the phase speed of a propagating wave, is the partial derivative of the pressure with respect to density (2). Here the gas is held at a constant temperature. That is why I jumped on Albert when he had the wrong variable held constant in his thermodynamic derivatives.

If we hold temperature constant, then this whole thing is a constant. If I take the derivative of pressure with respect to density, I can differentiate p with respect to ρ . It is, of course, 1. I can actually do a few other mathematical tricks. We use "a" in thermoacoustics as the sound speed rather than "c", because we have a lot of heat capacities floating around. The subscript "N" in (2), which is a Newtonian sound speed, is sometimes subscripted with a "T", to indicate that it is the isothermal sound speed.

It turns out that this was wrong because of the interaction of sound and heat. The correct answer is that a sound wave behaves like an adiabatic compression, not an isothermal compression. We take pV^γ as the adiabatic gas law, or if we are talking about a unit mass or a molar mass, then volume is inversely related to density and we can write the adiabatic equation of state for an ideal gas by putting $p\rho^{-\gamma}$ equal to some other constant, constant prime (4). We can differentiate that form of the expression to obtain (5).

That mathematical operation is a bit harder for me, so I take the natural logarithm. Then dp/p_m is equal to $\gamma dp/p$, so holding entropy constant gives us a value for the sound speed in (5). You end up with a result that is gamma times larger than the Newtonian sound speed, and that is the correct answer first obtained by Laplace.

If a wave propagates adiabatically, it is stiffer. This came up in Professor Sabatier's lectures. You modify the stiffness of a gas, depending on whether it is near a surface and held isothermal by thermal contact with the surface, or it is far from a surface and it cannot conduct its heat away so it must change its temperature adiabatically. That was the first problem.

[Transparency 5]

I will now prove why Laplace was right (far from solid surfaces) and why Newton was wrong (except near solid surfaces). This analysis will introduce another critical parameter in thermoacoustics. It is the distance that heat can diffuse in an acoustic cycle and it is known as the "thermal penetration depth."

Since you are all acousticians, I thought it might be useful to talk a little bit about another Newtonian phenomenon, which is heat transport. Newton wrote down the equivalent of Ohm's law for heat transport; the heat flux, q , shown in equation (7), is proportional to some transport coefficient known as the thermal conductivity, κ , times the temperature gradient, ∇T , in this case, the scalar field of temperature.

If we focus on only one dimension, we apply the equation to a chunk of material with a uniform cross-sectional area, A , and a thermal conductivity, κ , shown in the diagram, of length L . The temperature gradient is T_{hot} minus T_{cold} divided by the length of the sample. I am going to assume linear gradients in everything. We end up with something that looks just like Ohm's law for a resistor.

The resistance depends on the length and is inversely proportional to the area. There is a material constant here known as the electrical resistivity. We have exactly the same thing as for thermal conduction. This form is identical for any transport phenomenon, whether it be viscous shear, which happens to be a tensor, heat transport, which happens to be a vector, or electrical transport through any diffusive, time-irreversible phenomenon; it all is going to look like this.

These phenomena are described by linear response theory. I have written down here an equation (9) that Jim used just this morning. It is the shear force per unit area on a surface under a viscously flowing fluid. That force per unit area is related to the fluid velocity gradient by the

constant μ in equation (9). It is isomorphic to $V=IR$. This is a diffusion law. It is a consequence of irreversibility.

[Transparency 6]

We can take that law and we can do just what we do in acoustics when we look at plane waves or waves on a string or waves in a bar. We can create a differential form of that result by considering a thin slab of material, of unit cross-sectional area, that has some amount of heat, Q_{in} , entering it at a position x , and at a distance infinitesimally farther along, at $x+dx$, it has some amount of heat, Q_{out} , leaving the slab. Therefore, the net amount of heat deposited in this differential element, Q_{net} , is the difference of what went in minus what went out (10).

If an excess amount of heat goes in to this element, its temperature will increase. How much the temperature increases depends only on the heat capacity of the slab, which is the heat capacity per unit mass, c_p , times the mass per unit volume, ρ , times the volume of the slab. Therefore, the rate of temperature rise, $\partial T/\partial t$, times the heat capacity will depend on the net heat. The net heat is given by Newton's law of cooling, here in equation (11), in one-dimension, evaluated at x , minus (actually the minus turns into a plus) $\kappa \nabla T$, evaluated at $x+dx$.

We then use the second most powerful theorem in mathematics, the Taylor Series (12). What is the first most powerful theorem in mathematics? Substitution. Number two, Taylor Series. Number three, integration by parts, or the chain rule. I think you should take comfort, students, in the fact that I do not know number four and I am a tenured full professor with an endowment. You may not have to know as much mathematics as some people are telling you.

When you do that, the first term in the Taylor's expansion cancels the first term in (11), and we end up with the Laplacian in (13). From equation (11), if I bring the ρc_p over to the right-hand-side of (11), I form a quantity, χ , which has the dimensions of m^2/sec . Following Landau, I will define that quantity as a thermal diffusivity. Once again, we have an equation here that is the diffusion equation. It is not a wave equation.

Acousticians tend to work a lot with the wave equation. This is very important, also. Jay Maynard's talk was in-between. The wave equation has two time derivatives and two space derivatives. It is time reversal invariant. The diffusion equation has only time derivative but two space derivatives. The Schrodinger Equation has one time derivative but it is multiplied by $j = \sqrt{-1} = (-1)^{1/2}$. Another time derivative would generate an $j\omega$, so the Schrodinger Equation lives in this lovely no-man's land in-between the wave equation and diffusion equation, and that is why the "probability wave packets" spread; this is diffusive, this will spread.

[Transparency 7]

This diffusion equation shows up in all kinds of places. The Navier-Stokes equation, shown in (14), is a diffusion equation. I have allowed pressure gradients to accelerate the fluid. The acceleration of fluid is caused by pressure gradients but it is slowed by the diffusion of momentum through this equation, and there is a quantity μ/ρ that has units of length squared over time that is called the kinematic viscosity.

Fick's second law of diffusion describes a random walk. If I take a glass of water and a syringe with blue ink, and I shoot the blue ink into the water in the glass, there is a little blue cloud. Over time, it becomes a larger blue cloud. The time rate of change of the concentration of blue ink in the water depends on the gradients in the concentration; there is a diffusion of the constant, D , and its units, m^2/sec .

The same thing is true with Maxwell's equation in a conductor, a single time derivative, and two space derivatives.

This is a class of equations, and I point that out because I am going to solve only one and I am going to use the results from the others.

[Transparency 8]

As an acoustician, the first thing you do to solve an equation is you assume that the solution will look like a wave. I could solve any one of diffusion equations, but I am going to do the Fourier equation, the thermal diffusion equation, because temperature is scalar and makes the algebra a little easier.

What I would like you to assume is that I have a solid at mean temperature T_0 , in contact at this interface with a semi-infinite fluid that extends out to infinity in a direction $+y$. At this interface, the boundary condition is that the fluid temperature and the solid temperature have to be equal.

I am going to oscillate that solid temperature with an amplitude T_s (the temperature swings in the solid) at some frequency ω . I am going to assume that the spatial and temporal dependence of the resulting --

MR. CURRA: Excuse me, you said the temperature has to be equal. Should the heat flux be equal?

DR. GARRETT: Yes, there is energy continuity across the interface. I will argue that as you get closer and closer to the surface, the gradients get larger and larger if they are not equal, so there will always be sufficient "force" to make them equilibrate if you take the limit of an infinitesimal layer.

I am going to raise and lower the temperature sinusoidally in time. I am going to assume a wave-like solution and, again, as acousticians, we can take the time derivative if we remember the $\partial T_s / \partial t = j\omega T_s$. That gives me $j\omega$ on the left-hand-side of (17), and the two gradients gives me $-jk$ times $-jk$. The minus kills the minus, the $j \times j$ gives me another minus, and I get this minus of the diffusivity times T . We have got a perfectly linear equation.

I can now solve this equation for jk and I end up taking the square root of j , which is just $e^{j\pi/2}$, corresponding to a the vector going off at 45° . This produces a wavenumber, k , which has equal real and imaginary parts.

These wave-like solutions, these shear waves, these thermal waves, *et cetera*, are called evanescent waves. They have equal real and imaginary propagation vectors; this k in (18) has an equal real and imaginary part.

You need to remember that, because we commonly define a characteristic length (19), in this case called the thermal penetration depth, δ_k , which is a real number. It is the reciprocal of the modulus of the complex wave number, so it has the units of length. You can see its frequency dependence is given by the $\omega^{-1/2}$.

If we take this result for the complex wavenumber and plug it back into what we assumed the solution to be at the start, we find that the temperature variation, not the mean temperature, as a function of space and time in the fluid, has this form (20). It has an equal real and imaginary part. The distances are scaled by this characteristic thermal penetration length, δ_k . The temperature oscillations decay exponentially away from the plate with that same characteristic

length, and since it is a linear system, the thermal wave has the same frequency as the "forcing function."

[Transparency 9]

I can plot the real and imaginary parts for the case above (here I left it in complex notation rather than breaking it out into sine and cosine components). Tonight in the discussion session, I will show you some computer simulations where you can actually see the temperatures oscillating and you will get a better feel for the space-time dependence.

At this point, at the interface between the fluid and the solid, if the temperature of the plate is T_1 above ambient, then the temperature in the fluid adjacent to the plate directly is exactly T_1 , above ambient and the temperatures are in phase. The in-phase component is shown with a value of 1, the out-of-phase component, the 90° , or quadrature component, is 0. Then as we move out from the plate, the real part decays and the imaginary part reaches a maximum value and decays at greater distances. By the time the distance from the plate into the fluid exceeds three times the thermal penetration depth, δ_k , you are pretty much away from any influence of the plate.

The problem that you see in acoustics frequently is the inverse of the oscillating solid substrate temperature problem. You frequently encounter a substrate whose temperature is fixed and a fluid whose temperature is oscillating due to adiabatic sound waves in the fluid. The thermal penetration depth is the characteristic exponential distance over which the adiabatic variations would decay to zero at the interface. That is what I have shown in the lower illustration.

Far from the plate, the temperature of the liquid is oscillating between T_1 and $-T_1$, about ambient, and that oscillation amplitude decays exponentially as you go toward the plate. At the plate both the real and the imaginary parts of that temperature oscillation must vanish, because the plate has a much higher heat capacity than the gas.

This is true for most systems if you are talking about gases in contact with solids. I am going to restrict myself to gases, but a lot of very good work was done at Los Alamos with thermoacoustic engines that use liquid metals such as sodium-potassium eutectic, and that assumption of an overwhelming heat capacity in the solid is not good. There are not a lot of things that have high heat capacity compared to liquid sodium. Molybdenum was the best they could do, so the temperature of their stacks were being dragged up and down by the temperature oscillation due to the standing sound waves in the liquid metal.

This talk is titled "Thermoacoustics made simple," so I am going to leave it simple and not talk about liquid metal engines.

[Transparency 10]

I want to milk isomorphism between diffusion equations of different types so that I do not have to do more algebra. If you take a plate in a fluid and you oscillate the plate in a direction that is normal to its surface, you know what happens, you are all acousticians. It radiates sound.

What happens if you oscillate the plate in a direction that is parallel to its surface? If it is in a fluid with a non-zero viscosity, it drags the fluid close to the plate with the plate, because there is a nonslip boundary condition, and an evanescent wave comes out. If I rotate that, you will see that. The characteristic length over which the transverse fluid motion decays is given by the viscous penetration depth (23), δ_μ , as you would expect: the diffusion coefficient divided by the square root of the frequency, ω .

For air and for all of the 'fridges that Tom and I have built, the penetration depths, both viscous and thermal, are always about 100 μm . For air under normal conditions, the viscous penetration depth is 2.1 mm over the square root of frequency in Hertz, as shown under (23).

You can do the same thing with Maxwell's equation in a conducting medium. You get the same coefficient; you get a length known as the "skin depth" (usually it is not subscripted). In seawater, that length is roughly 30 m at 60 Hz, which is why most of Wisconsin is a very large radio antenna. It took all of Wisconsin to build a big enough antenna so they could communicate with submarines at a frequency low enough that the electromagnetic radiation would penetrate into seawater.

At about a megahertz there is a 2 cm penetration into water. In copper, there is 2 cm penetration at 1 kHz because copper is a much better conductor than sea water. The skin depth in copper is only a few microns in a megahertz. That is why waveguides generally work well.

The same result is obtained with mass diffusion, except it is called a diffusion length (25), but the people who are concerned about mass diffusion usually do not express their result in terms of frequency, they care about its reciprocal: time. The bottom line is you can exploit the isomorphism and save yourself a lot of algebra.

[Transparency 11]

This now allows us to answer the question of why Laplace was right and why Newton was wrong. Again, you remember equation (17). This is just without the T1's there, this is what we got from the dispersion law assuming a wave-like solution.

I can multiply both sides by ω , take the square root of ω/k , and find that the phase velocity for periodic thermal disturbances is dispersive. The phase speed is proportional to $\sqrt{\omega}$. One could see that easily by treating the wave number as being the reciprocal of the thermal penetration depth.

Therefore, since the speed of sound in a gas is constant (nondispersive) and the speed of heat is dispersive, there has to be some point at which these speeds are equal. That is the point at which pressure propagates at the same speed as heat. If you solve for that critical frequency (27), it comes out to be about a gigahertz in bulk, which means that Laplace is always right.

At a gigahertz the wavelength is less than the molecular mean-free-path, so the whole continuum assumption is invalidated. In gases, you never get to isothermal unless you choose to do that by putting the gas in a porous medium. It is commonly done in loudspeaker enclosures. The enclosures are stuffed with fiberglass. You want to make the volume of that box as large as you can; why would you ever put in something that uses up volume?

You put fiberglass in and it creates a matrix that has a high heat capacity. The gas compressions and expansions are isothermal rather than adiabatic and the stiffness of the gas is lowered by a factor of γ less than it would be if it were an empty box. Even though you are excluding volume, you are increasing the compressibility of the gas by making it isothermal rather than adiabatic. You have increased the effective volume by making the same volume less stiff. At high frequencies, the fiberglass also damps out standing wave resonances inside the enclosure.

The point is, this is the reason why Laplace was right and Newton was wrong. Sound is much faster than heat at ordinary frequencies, so if you get a compression at the wave crests that is hot, you get an expansion at the wave troughs that is cool. In the time that it has before the

situation reverses itself in the next half-cycle, not much heat can diffuse between the hotter crests and the colder troughs; it cannot get rid of that temperature difference by thermal conduction through the gas. A little of the heat diffuses and that is why you have thermal attenuation in sound waves -- the attenuation is proportional to ω^2 , and it is very tiny at audio frequencies. Thermal diffusion does not have time to wipe out that adiabatically-induced temperature difference.

DR. HAMILTON: At even higher frequencies, it should stay adiabatic.

DR. GARRETT: As you go to higher frequencies, the wavelength is getting shorter faster than the thermal diffusion length is getting shorter, because the wavelength is proportional to ω^{-1} , but the penetration depth is proportional to $\omega^{-1/2}$. As you go to higher frequencies, the wavelength is getting shorter faster for sound than it is for heat, but by the time they actually catch up, the wavelength and the thermal penetration depth are equal both about equal to the molecular mean-free-path. The continuum approximation has become invalid.

DR. HAMILTON: I never thought about it from that -- I like your wavelength description, that at low frequencies the peaks and troughs are so far apart that the heat cannot get from the positive to the negative.

[Transparency 12]

DR. GARRETT: I think what I will do is give you one more derivation here as background -- this is the last sort of "backgroundish" slide on fundamentals. Now that I have shown you that the wave propagation should be adiabatic, we know that there must be a temperature swing, because when we compress the gas, the heat cannot go out and it is going to raise the temperature of the gas.

If we have the adiabatic equation of state, just the repeat of equation (4), and we know the ideal gas law, we can derive the amplitude of the temperature swing in terms of the acoustically induced pressure oscillations. This derivation is actually more elegant than the one I learned; this approach is Tom Gabrielson's. We know that pV^γ is equal to a constant and we know that pV is equal to RT .

What Tom does is he takes pV and raises it to the gamma power. That is not what you want. But he then divides that out by p to the gamma and multiplies by p . We can substitute RT for pV , so now we have R , which is another constant, T^γ times $p^{1-\gamma}$, is a constant, as I have shown here in equation (30). If we take the natural log, again, and differentiate, we find that the change in pressure over the mean pressure is related to the change in temperature over the mean temperature, and in equation (32), we say that the magnitude of the temperature swing far from any surface in the adiabatic limit is $(\gamma-1)/\gamma$ times P_1/P_m (P_m means the mean pressure) times T_m .

I have written this in another way on the right-hand-side of (32). I will not derive this thermodynamic identity, but we can write coefficient of p_1 as $T\beta$ -- β is the thermal expansion coefficient at constant pressure -- over $\rho_m c_p$, the heat capacity at constant pressure -- times p_1 .

This is why we like inert gases for thermoacoustic devices. γ equals 5/3 for an inert gas. In fact, it is identically 5/3 at sufficiently low pressures. Mike (Moldover) is going to talk about this. The identity is strong enough that the Universal Gas Constant, R , is defined based on $\gamma = 5/3$, in inert gases.

We can look at the values for temperature swings given at the bottom of Slide 12. For normal speech, let's say its sound pressure level is 74 dB_{SPL}. That is a root-mean-square (RMS)

value that is one 10^{th} of a Pascal (0.1 Pa). For air, $\gamma = 1.4027$. That is an interesting number, because air is mostly nitrogen and oxygen. It is a little bit higher than what you get (7/5) if it were just diatomic N_2 and O_2 , because there is about 1% argon in there, and it is not as high as you would expect for the argon, because there is 334 ppm of CO_2 , which, of course, is triatomic. If you just add up the effects of the four most common components of air, you get that number.

Speech has, at the 74 dB_{SPL} level, 83 μK of RMS temperature variations. It will not cook a chicken. Let me tell you what those numbers are for an acoustic refrigerator I am going to talk about today. The SETAC refrigerator has pressure amplitudes of 65 kPa at full power, which is a large RMS amplitude. That is basically + or - 1 atm. The ambient pressure is 2.1 MP, roughly 20 atm. The gas is inert; it is a mixture of helium and argon.

There you start to see a substantial temperature change just due to the adiabatic compressions and expansions. If we look at its peak-to-peak temperature span in Fahrenheit, it is a big number. Is Hank (Bass) here? Yes. This is the difference between summer in Monterey and summer in Mississippi.

I did not want to come to the break without insulting somebody. You have got 10 minutes, then we are going to apply that to standing waves.

DR. GARRETT: Now that we know that there is a characteristic length, the thermal penetration depth, δ_K , over which heat can diffuse in a time approximately $T/2\pi = \omega^{-1}$, and we know the adiabatically-induced temperature variations in a sound wave, I would like to apply these results to derive what I consider to be some fascinating but very, very simple results. First, I would like to calculate a temperature gradient that would be characteristic in a standing wave of a gas undergoing adiabatic expansion and compression.

[Transparency 13]

To do that, I am going to assume that we have a tube that has a rigid termination at one end. At the other end, I have a piston moving with some volumetric velocity at a frequency ω . The tube has a length L . We will excite the tube with the piston at the fundamental, half-wavelength, standing wave mode. This means that the pressure are pressure maxima at $x = 0$ and $x = L$. The pressure has a cosine distribution along the tube.

If we use Euler's equation to solve for the velocity, I can solve, again using Fourier decomposition, the time derivative of the acoustic part of the velocity. I am assuming now that there is no mean gas flow through the tube so u_0 is zero. The particle velocity amplitude, u_1 , is related to the pressure gradient from Euler's equation (35), so since we know the pressure, we can write down the pressure gradient and solve for u_1 , as shown in (36).

There is a 90° time-phase shift, a j , between the pressure and the velocity. If I solve, again, for the displacement, the integral can be obtained by division by $j\omega$. I integrate the velocity, and come up with an expression for the particle displacement (37), so at any given position x in the tube, a gas particle will be driven by an acoustic wave of amplitude of $P_A/\gamma P_m$ which this is just the acoustic Mach number. $\lambda/2\pi$ is called the reduced wavelength, k^{-1} . At the center of the tube, where the velocity is the maximum, the displacements are on the order of the reduced wavelength times the acoustic Mach number.

[Transparency 14]

I have a picture here now of a parcel of gas at some location x within that standing wave tube, but I have not specified what that location is. The sound wave is driving that parcel to the left of that point and to the right of that point by a peak displacement amplitude x_1 .

At the right-most location, $x + x_1$, the gas parcel temperature is reduced from its mean value by that adiabatic amount, T_1 . At the left-most location, it is increased by that adiabatic amplitude. We have already written this equation (38). I have now substituted a standing wave pressure field for p_1 . The particle position is given by my last equation on the previous viewgraph (37) and, therefore, there is an acoustically-induced temperature gradient (40) associated with a standing wave sound field.

When the gas is on the right, it is colder. It is hotter over at the left. That is true at all times; therefore, there is a static temperature gradient. Twice T_1 is the ΔT and twice x_1 is the Δx . Again, I am assuming linear gradients all the time since I have also assumed that there is no mean flow of the gas.

We find that the ratio of those quantities is a number that is given by the product of something that is characteristic of the gas, $\gamma - 1$, the polytropic coefficient minus 1; the mean temperature of the gas, T_m , scaled by the wave number, k ; times the cotangent of the position. The pressure amplitude cancels out! This is a linear result. It does not matter how hard you drive the gas, that pressure gradient is a number and that number diverges as you approach the rigid end where the pressures are large, the temperatures are large, but you are up against a "brick wall," so u_1 is zero. Anything divided by zero is going to be big!

This is a very conservative set of numbers below (40). (Tom will back me up on this). A stack mean location times the wavenumber is maybe a quarter and the wavelength is a couple of meters -- Tom's wavelength in the demo is four times the length of his resonator, maybe that is a meter in his case, bigger 'fridges may be a couple of meters. We will take the mean temperature as somewhere between room temperature and a good chill.

This temperature gradient turns out to be 22°C/cm . That means in 2 inches, 5 cm, in principle, you can turn steam into ice. These are not small temperature spans. At the surface, of course, you can write T_1 down, because the surface is assumed to have a high heat capacity. T_1 at $y=0$, right on the surface is zero. So it goes from a fairly large adiabatic T_1 in the gas to T_m and the distance over which that temperature decays is the thermal penetration depth, δ_k , which is why I spent so much time deriving it.

[Transparency 15]

We now get to the one and only important viewgraph and we will spend some time on it: the moving parcel picture, as it is known amongst the cognoscenti. This is quite a bit different from what Tom showed you, but you can build one that is this simple.

It has a resonator -- this is just a tube of uniform cross-section. I will show you a stack and a heat exchanger, but Tom has shown you a prettier stack (I will show you an uglier heat exchanger). This loudspeaker could be just an ordinary commercially available loudspeaker. We drive the loudspeaker at a frequency that will excite gas in the resonator to oscillate in the fundamental, half-wavelength standing wave mode. What does that mean?

That means that the gas near the center of the resonator is sloshing back and forth. The gas near both ends acts as gas springs. The gas at the ends is compressed at one end while it is rarefied at the other end. One-half cycle later, the situation is reversed. The pressure on the

central gas slug due to the gas springs forces it back to the center. That gas slug also has inertia, there is a mass associated with that gas, so it overshoots, and you have got this gas oscillating back and forth almost like a mass contained between two springs.

At the center of the tube, as I showed before, you have only velocity, no pressure swings. At the ends, you have only pressure swings no velocity. At any intermediate position, you will have some mix of pressure swings and velocity.

The stack is chosen to have that tight plate-to-plate spacing. Typically in acoustic refrigerators, the spacing is roughly equal to two thermal penetration depths. This means that the stack acts as a porous medium in which all the gas can be in modestly good thermal contact with either the roof or the floor in the amount of time that that is available during an acoustic cycle; roughly a couple of milliseconds.

What I have shown in the lower portion of the slide is I have "zoomed in" on a little piece of stack. I have colored a million of these indistinguishable particles red. I am going to talk you through an acoustic cycle for three different cases: in the first case, the substrate has a spatially uniform temperature. Next, we will repeat the cycle when the stack has the critical temperature gradient imposed on it. Finally, we will look at the cycle in the presence of a temperature gradient that exceeds the critical temperature gradient. Hopefully, this will explain to you, at least qualitatively, what makes this process a refrigerator and what makes it an acoustic motor. This should encompass both of Dr. Hofler's demonstrations. I will do it in more detail and derive an equation for the heat transport shortly thereafter.

As you know, every parcel of gas is undergoing a sinusoidal oscillation. If you remember Jay's lecture, the wave function is peaked at the extremes. Why? Because the gas comes to rest, changes direction, zooms through the center. Basically, it stops at one end, and that is why I am drawing it as an articulated cycle: a quick adiabatic compression to the left; a rest; then a quick adiabatic expansion to the right and another rest. It is not exactly sinusoidal, but there will be bigger swindles than that before this picture is complete.

We start out with a uniform temperature, say room temperature T_0 on this section of stack (substrate). The gas is within this thermal penetration depth and, therefore, if I am starting here, it has had plenty of time to equilibrate to the stack temperature. Everything is at equilibrium; everything is at T_0 .

The sound wave compresses the gas, moving it quickly to the left. The adiabatic compression raises its temperature by two units. Let's call this elevated temperature T^{++} . The gas then comes to rest. It is within the thermal penetration distance, δ_k , and it is hotter than the substrate that is still at T_0 . Newton's law of cooling says heat is going to flow from hot to cold, so during that "rest" period, the temperature of the gas is going to be reduced while the temperature of substrate will be increased slightly. The gas parcel will cool when it deposits this heat, Q_{hot} , to the substrate from a temperature T^{++} to a temperature T^+ .

The third quarter cycle will be an adiabatic expansion that was identical to the compression. Since the gas parcel started its expansion from the lower temperature, T^+ , it must end its expansion at a lower temperature, T . It is now at its right-most location. I have drawn the gas parcel above its original position here only for clarity, so you could see it -- it is actually still within the thermal penetration depth. It is, once again, approximately at rest. It is colder than the

substrate and sucks an amount of heat Q_{cold} out of the substrate and returns to its starting temperature, T_0 .

What we have illustrated here is a classic refrigeration cycle. We have done work in the form of sound and we have pumped heat up against the direction it would flow due to thermal diffusion. That is what Tom was doing in his first demo. If we zoom back to the figure at the top of the slide, what we see is the particle displacement is about a millimeter. The total length of the stack is about 2 cm, so there is a bucket brigade of 20 of these parcels handing off the heat from one end of the stack to the other.

The right-most parcel hops off the end, picks up some heat from the cold heat exchanger, dances it up to successively higher temperatures, and dumps it on the hot heat exchanger at the left of the stack. As you saw in Tom's demo, that hot heat exchanger is thermally connected to a fan and it is big. It is not going to heat up much. Therefore, the temperature of the cold heat exchanger is going to drop like a rock. That is what you saw in Tom's first demonstration.

This is extraordinarily important in terms of physics and John Wheatley recognized this; he is one of the founding fathers. Yes?

MR. HODGSON: How far between steps 1 and 3?

DR. GARRETT: About a millimeter. In Tom's big prime mover or Greg's prime mover, the gas in the center is moving by a substantial fraction of a meter. In Tom's demo 'fridge, the gas is probably moving a little less than a millimeter, probably about half a millimeter in your 'fridge. The gas in the stack of the Hofler tube is moving probably close to a couple of millimeters.

What is important in the physics of this is the thermoacoustic process is an intrinsically irreversible engine in the sense that we are counting on Newton's law of cooling, an irreversible phenomenon, to give us the proper phasing between compressions, expansions and heat transfers. If you look at an engine with which you have probably more familiarity, an automobile engine, say, it may require as many as 24 valves to provide the proper phasing between intake, exhaust, compression and expansion.

Each one of those valves is articulated by a rocker arm. Each one of those rockers has a push rod. Each one of those push rods is sitting on a cam. There are an awful lot of parts in that engine whose sole purpose is to give you the proper phasing between compressions and expansions, exhaust and intake, that are required to make that engine run.

The thermoacoustic engine had no parts to provide the phasing. It was phased naturally. We call this a natural engine. It means the prime movers will run on Granola. (Laughter)

This is an important feature of thermoacoustics. It allows construction of a very, very simple engine. The phasing is being done by natural processes, not by mechanical contrivances, and that is one of the great strengths of thermoacoustics: it can provide the proper phasing without having a lot of mechanical linkages.

PARTICIPANT: Doesn't that make it less efficient?

DR. GARRETT: That may be true. You are going to get a lot of those "may be true" answers, just like the answer about the 7° taper angle in the tube to suppress turbulence in oscillatory flow may be true. That does not stop us from building them. We do not know. We really care. It is not as if we do not know and we do not care. We really care. We do not know.

Let me continue with this. Is everybody happy with this refrigeration process and its utter simplicity? You can see that eventually, if there were no dissipation, the temperature of the left end of the stack near the hot heat exchanger would come up to T^{**} , and the right-hand-side would drop down to T^* . At that point, the entire heat transport process would come to a halt. If the gas is at T^* , and the substrate is at T^* , there is no temperature difference across which the heat will be driven into the gas or out of the gas and onto the substrate. That point corresponds to the critical temperature gradient.

When the gas parcel sees its own adiabatic temperature lapse mirrored on the substrate, that is the line of demarcation between refrigerators (where you put in work and move heat up a temperature gradient) and prime movers or motors (where you put in heat and generate work in the form of sound).

There is nothing that prevents you from taking a torch to one end, or applying liquid nitrogen to the other end, as Tom did in his Hofler tube demo. I could force one end to be T quintuple minus and this to be T triple plus and intentionally exceed that critical temperature gradient. If that happens, then when the gas moves to the left, it is compressed, and heat will diffuse into the parcel from the still hotter substrate. The incoming heat will do work on the gas parcel at high pressures. When it is expanded, heat will come out of it. The net effect in going around the entire pV (indicator) diagram is you are putting work into the gas from the heat flow out of the stack.

Let me review that. Assume we have a substrate whose temperature is uniform or it has a temperature gradient that is less than the critical temperature gradient. We impose a sound wave with a loudspeaker and the sound will pump heat the wrong way, the way heat does not like to flow naturally. When we reach the critical temperature gradient the heat pumping halts. The critical temperature gradient is the line of demarcation between the refrigeration regime and the prime mover regime. If we include other forms of power dissipation such as viscous scrubbing on the stack and thermal conduction along the stack and through the gas, it is not a thin line but more like a demilitarized zone. There is some region of temperature gradient conditions that will give neither a refrigerator nor a prime mover.

I have not included thermal conduction back along the gas and in the stack and I have not included viscosity, but if I did, you could not get to the critical temperature gradient. That is an idealized limit. If you could get to it, you would obtain Carnot efficiency. The thermoacoustic device will have the same efficiency as any other heat engine if you pull no power out of it.

On the other hand, if I were to make this T^{***} and leave this at T^* , that may not be a sufficient temperature gradient to generate enough power so that this thing will go through its Hopf bifurcation, if we use Prof. Laterbourn's terminology. Although we call it a critical temperature gradient, although we have defined it in the absence of viscosity and that does not apply to real systems, it is such a useful concept that we retain it, even in the most complicated theories, because it is intuitively very powerful, as you will see.

I will point out, also, at this stage that we do have heat transport between 2 different temperatures and, therefore, there is entropy generation. This is the thermal dissipation that you have near a surface. Any time there is heat transport across a temperature gradient, there is entropy generation and that is irreversible, so there is loss, even in the absence of viscosity.

If there were viscosity, of course, when you are scrubbing along here, you are generating loss, and I am going to talk about that, but I wanted to point it out while that picture was on the screen.

[Transparency 16]

I would now like to milk that picture a little bit harder and calculate what that heat transfer is in a very simple case. Let's say that I look at one half-cycle. I treat the gas between x and $x-x_1$ as being a chunk of gas that is going to get moved a distance x_1 . It is going to be colder than the substrate. I am going to assume no temperature gradient in the gas or the substrate; that is, uniform temperature along the stack plate.

I am going to let the gas parcel cool adiabatically. Being colder than the substrate, it will suck Q_{cold} into the gas, move over to the other position, heat up adiabatically, and dump an amount of heat Q_{hot} at the other position. I am going to calculate the heat transport in a very crude manner. My justification for doing it crudely is that the final result will about right. It will give you the right dependencies on pressure amplitude and frequency. We have good models that can calculate the heat transport correctly in very great detail, including the viscosity and including the heat conduction, and accommodating stacks of different pore shapes.

What happens? This gas parcel is at temperature $T_{\text{mean}} - T_1$ over here, it is at T_{mean} in the center, so on the average, the temperature of this whole thing, let's say, is $T_1/2$ lower than the mean -- just average everything.

For historical reasons we treat the width of the plate as being the plate perimeter Π , divided by two, because the same heat pumping process is happening on the bottom surface of the plate.

We will assume that the gas that can actively pump heat is within a single thermal penetration depth, δ_k , from the surface of the plate, so that we have a well-defined volume of gas. The amount of heat that is going to come into this volume is going to be the heat capacity of that slab, x_1 long, $\Pi/2$ wide, and δ_k high; times the temperature difference between the substrate and the plate. The heat capacity of this gas parcel is the heat capacity per unit mass times, c_p , times the mass/unit volume, ρ , times the volume.

When the gas moves to the left, it will be compressed, and the gas here will be at a temperature T_1 above T_{mean} (of course, it will be T_{mean} at the center position, x). I will split the difference and call the average temperature difference $T_{\text{mean}} + T_1/2$. It has the same size as before; so, the net amount of heat that has been moved from is the difference of those two (43). Remember, the leading terms in (41) and (42) will cancel when I subtract them. I get two $x_1/2$, which is x_1 . I get the thermally active cross-sectional area of this face, $\Pi\delta_k/4$, and the heat capacity per unit volume, ρc_p .

Generally, when you do acoustics, you do not measure or calculate displacements; you use velocity. I am going to express, again, being very fast and loose with my imaginary quantities, x_1 as u_1/ω and T_1 the pressure given by the adiabatic gas law. Eq. (44) represents the amount of heat transported by the sound wave to the left by the designated parcel during one-half of the acoustic cycle.

[Transparency 17]

I have to time-differentiate the heat in (44) to get the power, which means multiplying by ω again, so the u_1/ω gets canceled by the ω in equation (45). I then obtain an expression that is the

product of the mean temperature of the gas times the thermal expansion coefficient. For an ideal gas, $T\beta=1$.

The effective cross-sectional area for the heat transport is the perimeter over two times the penetration depth. The product of pressure times velocity is the acoustic intensity. Eq. (45) is valid when we assumed that there was no static temperature gradient on the stack plate.

I am going to modify (45) for conditions where there is a temperature gradient in the plate and you can see that if I define capital gamma, Γ , which is the ratio of the mean temperature gradient in the gas and on the stack to critical gradient. It is easy to see in (46) that when we are at critical gradient; no heat would be transported, as we agreed in the moving parcel picture.

I am going to multiply the heat-transport rate (45), by Γ^{-1} . When there is no temperature gradient on the plate, I get the maximum amount of heat transport. When there is a critical gradient on the plate, the heat transport stops. In fact, when Γ is much larger than 1, the heat goes the other way. We are going to do that with prime movers, our first example.

If you do the thermo-hydrodynamical calculation correctly, what you want to do is calculate the total energy flux, which for these conditions would be the enthalpy flux. It would include both the potential energy and the thermal energy. It does not include the kinetic energy for Mark Hamilton's reasons, because the kinetic energy is intrinsically second order and its transport is u_1^2 times u_1 . It is third order and, as Landau said, we can throw those out without degrading the solution at this level.

When you do this properly, (47) is the enthalpy flux. It turns out for an ideal gas the term which includes the acoustic intensity actually vanishes since $T\beta=1$, but there are other working fluids than ideal gases. The calculation I did, even though I said it was only the heat transport, was, for an ideal gas, also the total energy.

I am going to go one step further. After doing this correctly you obtain the same result in the absence of viscosity and we are saying that energy is transported in one direction if the temperature gradient is less than critical and in another direction (actually, the negative direction) if it is super-critical -- I have added on the thermal conduction through the gas in (48). The half-spacing between the surfaces of the stack plates is y_0 , κ is the gas thermal conductivity. l is the half-thickness of the substrate and thermal conductivity of the substrate is κ_s , so that last term in (48) is the heat transport due to Newton's law of cooling through the gas and the stack. I have thrown that in as long as I am doing the total energy, and that is where I am going to quit throwing things in.

We can ask ourselves now what is the efficiency of this process. The problem is we cannot answer that question yet, because we have not calculated the dissipation. That is what I am going to do on my next slide.

Before I do that, again, I want to point out, whichever equation you use, that this heat transport is second order. That is what the subscript "2" is, just like Mark's lectures. It is proportional to linear combinations of two first-order quantities. If all we wanted to do was transport heat in a standing wave, you would put the stack at a quarter wavelength from the rigid end because the pressure is large and the velocity is large.

You can see that if you are building a refrigerator, if you are designing a real device, there is already an ugly little trade-off that you are going to have to face. The temperature gradient

gets bigger and bigger for a fixed wavelength in the gas, as you go closer to the rigid end, where the velocity vanishes. Remember, $\cot kx$ in (40) diverges as $kx \rightarrow 0$.

If you want to span a certain temperature difference, you can have a very short stack very close to the rigid end, because you have a very high critical gradient. On the other hand, if you do that, you are not going to pump much heat, because the thing that caused your gradient to get large was u_1 vanishing. Eq. (46) shows that the heat transport is proportional to p_1 times u_1 .

It is going to require an engineering trade-off. I have a quote from John Wheatley and maybe this is the time to put it up here: "Heat engines are a compromise between the crisp ideals discussed in thermodynamic textbooks and the clanking, hissing realities of irreversible processes. This compromise produces wonderful machines such as the automobile engine and the household refrigerator. In designing the real devices the goal is not to approach thermodynamic ideals by reducing irreversibilities but to balance cost, efficiency, size, power, reliability, simplicity, and other factors important to the needs of a particular application."

Prof. Wheatley's point-of-view is very important. I probably will use this again, but I think this is a good time to bring it up. There is going to be a tradeoff between heat-pumping power and efficiency. There is going to be a tradeoff between temperature span for a given size of stack and heat-pumping capabilities. Live with it. That is the way of the real world.

[Transparency 18]

We have to determine what the losses are in order to calculate what the overall efficiency of the process might be -- under any conditions. Since we have surfaces, we have thermoviscous dissipation -- scrubbing -- and we have irreversible heat conduction.

That gives rise to two terms that probably are familiar to most of you. This is what is known as Stokes-Kirchhoff attenuation, or boundary layer attenuation, in ducts or resonators. Lower-case $e\text{-dot}$ is the energy dissipation per unit area. It is equal to a term that looks like the kinetic energy density times the viscous penetration depth. Multiplication by the frequency gives you a rate. The other term in (49) depends on the polytropic coefficient, γ , and the potential energy or the pressure squared, in the gas. That is the thermal part.

You can see that if the polytropic coefficient is 1, then you have isothermal behavior, there is no ΔT in the gas. If there is no ΔT , there is no irreversible heat conduction and this term should go away, and it does go away. The viscous term, of course, does not go away; there is always friction.

The viscous term is simple. The thermal term is important. Let me derive the viscous term for you, just because it is simple and because I would like you to believe that this is something that would make sense. Power is force times velocity (50). The force per unit area for viscosity is that shear stress, the Newtonian shear stress.

Their product is u_1 times $\mu \partial u_1 / \partial y$. This term in equation (51) is just that product. But again, we are acousticians and we define our characteristic lengths in terms of exponentials. Mechanical engineers do not do that; they define the characteristic length for shear flows as the distance between the wall and the point at which the velocity equals 99% of the free stream velocity. E.E.'s do not always use exponential lengths either. They define the rise time from 10% to 90%. That leads to a lot of ugly little numbers.

Acousticians always use potential and, therefore, we take derivatives by just dividing by the exponential length. This goes directly to μ times u_1 , time u_1 over the viscous penetration depth,

which gives you that gradient. The result is divided by two because we are time-averaging and this is a quadratic quantity. This is this "RMS 2" in (52).

The surface area is the length of the plate, Δx , times the perimeter of the plate, Π . The viscous penetration depth can be used to substitute for the μ and you end up with exactly the expression that you have on the first term in (49). You have got the kinetic density energy times the frequency times an area times a viscous penetration depth. There is nothing magical, it is a very simple thing to derive on the back of an envelope.

DR. MOLDOVER: You have plug flow and a laminar boundary layer. Do you get two amplitudes where turbulence in the boundary layer is important?

DR. GARRETT: We hope not. I am going to show you some data where we see deviation. There has been research done. Merkli and Thoman defined a boundary-layer Reynolds number and in my early days of 'fridge designing, I used to use set the acoustic Reynolds numbers to less than 300 as being the maximum amplitude. I would try to do a design, when I was trading off stack position and amplitude, keeping well below that limit.

There is a lot of research people are doing, using laser-Doppler velocimeters, to look at what is happening there. The fact is, there is so much in-between the plate and the bulk, or the heat exchanger and the bulk, there is so much turbulence that nobody has ever bothered to get a good picture of what is going on in the plates. Or, I am wrong.

DR. HOFER: Where there are small tubes, you get quite a large boost in the velocity. We were talking about flares at the ends of the tubes, that is one high-amplitude issue. The other one is following the peak velocity -- actually, the velocity peaks just before you get to the flare, you can fairly high "boundary-layer" Reynolds numbers there.

DR. GARRETT: I thought Mike was referring specifically to the stack.

DR. HOFER: Yes, certainly, the stack would be the most critical question, but that is a lot lower velocity. Usually where you would probably see it first, if you are using this resonator with the small tube, is in the resonator tube but, hopefully, that is the weaker effect than if it happened in the stack.

Hopefully we are avoiding it in the stack, although whenever you get to the end of a heat-exchanger fin or a stack plate you are probably going to get vortex shedding there before you get turbulence, real boundary turbulence.

[Transparency 19]

DR. GARRETT: Let me calculate in a little bit more detail, just to show you how one might calculate thermal dissipation. What you need to do is calculate the power, $d(pV)/dt$ which is just pressure over density times dp/dt (equation (54)). We can expand the total derivative of density in terms of the time derivative and the convective derivative (55). Expanding p to first order in terms of the mean pressure, p_m , and the deviation, p_i , we end up with a product which has four terms.

If we time-average, three of those four terms disappear. You end up with a term that is $j p_i p_i$, and you want to calculate the time average. Complex "j" gives you zero, the power factor becomes zero, because the phase is 90° . But there is this additional lag. Remember that the temperature had a real and an imaginary part and it turns out that the part that is left over is the imaginary part since its product with j is real.

The power-per-unit volume that is absorbed due to the expansions and contractions of the gas in its thermal motion is a quantity that is proportional to the amplitude and proportional to the imaginary part of the temperature, the part that kills the j when you do the time averaging.

If you integrate in the direction normal to the plate surface, you come up with this result that I added for the thermal dissipation (57). All of it was well known in 1957 except this $\Gamma-1$. This is what was called Kirchhoff attenuation.

Of course, this $\Gamma-1$ makes perfect sense, because if you reach critical gradient, the thermal dissipation disappears and there had better not be any irreversible conduction. If γ goes to 1 the loss disappears, because there would be no temperature swing in the first place.

What makes (57) important is that there is no limit on what the value of Γ can be. I can make the mean temperature gradient much bigger than the critical temperature gradient and, therefore, (57) can be producing power instead of dissipating power. That is, in some sense, what the Hofler tube demonstrated. We imposed a temperature gradient big enough that it could compensate for the losses that were intrinsically viscous.

[Transparency 20]

I am going to go through a review of all the fundamentals we covered, let you have a little break, and then we are going to look at this applied to several engines for which I have a great deal of narcissistic affection.

Sound waves in gases far from solid boundaries are adiabatic. That is because the wavelength of sound, λ , is always larger than this thermal penetration depth, δ_K . The thermal penetration depth tends to be, for ordinary frequencies, on the order of a 10^{th} of a millimeter, or 100 μm , wavelengths at those frequencies tend to be on the order of meters, and the heat cannot get from point A to point B before point B becomes point A and point A becomes point B.

If we take this result from the adiabatic gas law, which we say is correct far from surfaces, then we can express the temperature swings in terms of the pressure swings. If we apply this result to standing waves -- forgetting viscosity, just a simple standing wave -- we find that there is a critical temperature gradient that this adiabatic process generates, there is a temperature gradient. This we call the critical gradient. It depends on the wavelength of the sound but not the amplitude.

You know what the temperature gradient going to be in acoustics. It is going to be the mean temperature over the wavelength. It is going to be proportional to the same $\gamma-1$ factor that we had in (32) and it is going to depend on the ratio of the pressure that gives you the temperature swings to the velocity that gives you the displacements. That ratio will diverge at the rigid ends of a resonator (40).

This is true in the bulk, but at a solid substrate with a high heat capacity, it is not true; the temperature variations about the mean temperature must vanish. The distance over which the temperature can change is given by this characteristic length called the thermal penetration depth, δ_K .

The derivations I have done so far are what are called the short-stack approximation. I have assumed that the total length of the stack, Δx , is very small compared to the wavelength. The velocity of the gas at any point along the stack was basically the same, the magnitude of the pressure swings were basically constant.

We came up with an energy transport equation (48). The enthalpy transport, which is total energy, depends on the product of the pressure and the velocity and this $\Gamma-1$, which takes into account the mean temperature gradient on the stack. There is a diffusive part that you have just because heat flows from hot to cold; it is equal to the mean temperature gradient times the thermal conduction.

We have calculated the energy dissipated when we put that standing wave pressure and velocity distribution in this "short stack" and we have said that the dissipation has two terms. Both terms are positive definite in the sense that u_1^2 is going to be a positive number and p_1^2 is going to be positive.

The last term in (58) is always negative. Viscosity always absorbs energy due to friction. The "thermal" term in (58) is also negative if the temperature gradient is below critical. The irreversible thermal conduction eats up power and generates entropy. The viscous scrubbing eats up power and generates entropy. If we allow the gradient on the plate to exceed that critical temperature gradient, in which case the thermal term changes sign; it becomes positive, which is opposite from what we expect for a dissipative process.

We expect that for some value of Γ greater than unity, the two terms in (58) will be exactly equal this and we will have an instability. The resistance, the loss term, which is usually a positive number, will become negative. It will become a negative resistance oscillator and any pressure fluctuations will grow spontaneously.

After the break, I will apply this result to a solar version of the Hofler tube. Given those results, we will show how much heat can be transported. I would like to conclude this section by saying that everything I have talked about is occurring at a surface. This is very superficial science, thermoacoustics. I think that after the break you will see that we can use these concepts and these models to produce some very interesting devices.

[Transparency 21]

I have reproduced the work equation (59). This would be the energy dissipation equation if ignore the $\Gamma-1$. I would like to address the prime mover, which Prof. Laterbourn likes to call the Hopf instability. The losses will become negative and the system will initially exhibit exponential growth.

If we look at (59), we know that the viscous term is intrinsically negative. The magnitude of the two terms can be numerically equal for a sufficiently large temperature gradient. We can set these two terms equal and solve for what $(\Gamma-1)_{\text{onset}}$. That is called the onset condition.

We get a very simple result (60). The Prandtl Number is just the ratio of the viscous dissipation to the thermal dissipation. It is the ratio of the viscous penetration depth squared to the thermal penetration depth squared. The Prandtl Number, σ , is the property of the gas. It is why I made a mistake in Jim Sabatier's lecture this afternoon. I am used to inert gases where σ is always 2/3. It is 3/4 in air.

You can argue these are the short stack results, so this does not take into account the resonator, it does not take into account the radiation, which we admit is non-negligible in the Hofler tube case. But all of those other forms of dissipation: the resonator losses and the radiation losses, are all dependent on p_1^2 or u_1^2 . We can express them all in terms of p or u , as I have done here, to get rid of the p 's and the u 's -- that is what gave me the $\tan kx$ term.

If I included some more loss terms, I would still have the same form for $(G-1)_{\text{onset}}$; additional dissipative mechanisms would just require a little larger Γ to compensate for those losses, but they are all quadratic in the linear acoustic variables. They are all positive definite dissipation mechanisms, whether they be radiation, or scrubbing in the resonator, bad solder joints, or leaky stuff. It will go into onset somewhere unless it melts or catches fire first.

As in the case with most instability problems, the calculation of the onset (break-even) point is fairly simple. The calculation of the steady state point is very difficult, because some nonlinear process will limit the exponential growth. All the onset calculation tells you is that the linear solution diverges at this point. Initially, it is going to grow exponentially, but it is not going to grow to infinity, it is going to be limited by something. Either some nonlinear effect or you reach the point where you have used all of the available energy.

There is an onset point. Above onset what happens to that additional heat in a prime mover? If I put in more heat than it takes, does the onset point go up or does something else happen? The answer is not obvious. We have to put in a certain amount of heat to generate the onset temperature gradient: Let us call that the onset power. If we put in more power than that, it turns out that $(\Gamma-1)_{\text{onset}}$ remains constant and the acoustic amplitude grows.

If this simple model is correct, I can solve for the steady-state acoustic amplitude. Eq. (62) shows the pressure amplitude squared, at the rigid end in a resonator of uniform cross-section (which, of course, nobody uses in his right mind). (62) includes a whole bunch of trig, some material properties, and the difference between the heat that you put into the hot side, Q_{hot} , and the heat that it took to go unstable, Q_{onset} .

If I plot the temperature gradient, or equivalently, the ΔT , across the stack, it should rise to some level and then stay constant irrespective of the amount of additional heat supply. If I plot the acoustic amplitude squared versus the heat input to the hot side, there should be some non-zero intercept, which would correspond to the heat required to bring the engine to onset. At higher input powers, p^2/Q_{hot} should be a straight line.

That is the prediction of this simple theory. I am going to show you my favorite prime mover that substantiates that result experimentally.

[Transparency 22]

Along with this favorite prime mover, I am going to show you a large subset of my favorite scientists. This photograph was taken out at Penn State on the second floor patio. Matt Poesse is missing -- I think you shot this, Matt, so unfortunately he is not in the photo. On the left is G.W. Swift, the great Dr. Swift. Next to him is his sidekick, a Penn State graduate, Bill Ward. Next to Bill is Bob Smith, known as Mama Bob in the group. He is our engineer-technician in thermoacoustics. Next to him is a very tall gentleman named John Corey. He has a company called Clever Fellows Innovative Consortium. Most people would not have the guts to call their company Clever Fellows. He does and he gets away with it, he is that smart (I will show you some of his work later).

That is the evil Darth Garrett and his colleague, the evil Tom Gabrielson. Reh-lin Chen is down here praying that this thing will go into oscillation. This is a solar prime mover -- I will show you some better shots of it. Anat Grant, another of my graduate students, is working on the SETAC 'fridge that I will talk to you about later. Right next to the Fresnel lens is Robert M. Keolian, the dean of demos.

What they are posing with is a heat-driven thermoacoustic engine that contains a stack that is made out of the ceramic material that is used in a catalytic converter. This is a little trick that Hank Bass and Rich Raspet taught me. You can purchase this ceramic stack material from Corning. It has fairly low thermal conductivity and it is very inexpensive before they coat it with platinum to make it into a catalytic converter used in automobile exhaust systems.

On the back of the resonator is a beaker, an ordinary 500 ml beaker that has been cut off, to make a window. On the other side of the ceramic stack is a heat exchanger. That heat exchanger is soldered into a plate that has two turns of copper tube attached through which water will flow. Reh-lin has put some nice color photographs of the entire system on our thermoacoustics web site.

The tubes of water are connected to a bunch of heat exchange fins. We measure water flow rate with a turbine flow meter. A thermopile made from 10 type-T copper-constant thermocouple pairs, measures the differential temperature between the water inlet and outlet at the heat exchanger. If you know the density and heat capacity of the liquid, measure the flow rate and ΔT , you can calculate how much $Q = (dm/dt)c_p\Delta T$, is leaving the system where dm/dt is the mass flow rate of the water. It is a very accurate, very useful and convenient way to measure the heat that is coming out of the cold heat exchanger.

If you know how much work is being done acoustically, you know how much solar energy is being deposited onto the hot side, since the First Law of Thermodynamics requires that $Q_{\text{hot}} = Q_{\text{cold}} + W$.

[Transparency 23]

Here are data from that solar engine. The vertical axis is the stack temperature difference, so this is proportional to the temperature gradient, because the stack length is fixed. The horizontal axis is the electrical input power. These are not solar data. These data were taken with an electrical heater, because it is easier to measure and control.

The diamonds are pre-onset data. Above about 100 W, it breaks into oscillation and those points are shown on the graph as asterisks. You can see that if you increase the heater power by a factor of three above the onset power, the temperature gradient remains relatively constant.

On the lower graph, the acoustic amplitude vs. heat power input is plotted. The star on the lower graph is the only actual solar measurement; the triangles are electrical heater measurements. You can see that there is a nice straight line if we plot the acoustic pressure amplitude squared versus the input heat, so it really does work the way it is supposed to (63). The absolute numbers are correct, too, if you take into account the heat loss by radiation and by convection on that beaker.

[Transparency 24]

As Tom pointed out in one of his slides, there is no reason, once you have a prime mover, why you cannot slap another stack in the resonator and make a heat-driven 'fridge. so we have a ceramic stack. Here the window is a hemispherical glass shell. This fridge is designed to work at an internal pressure of 7 atm. This window is cute; it is made by Benthos. Benthos is a company that makes underwater instrumentation and they build glass pressure vessels that will handle 9000 psi. They will sell you half of one. They are quite nice. That is the window and it is big window because you do not want the concentrated sunlight hitting anything but the end of

that stack. You want everything else clear of that solar stream. It is a very dangerous experiment.

It says hot heat exchanger here, but there is no hot heat exchanger. There is an electrical heater, again, because of the meteorological problems in Pennsylvania. There is an intermediate exchanger at ambient temperature and then a refrigeration exchanger. This is designed to create ice in Africa. It is for storing medical supplies where there is no electricity.

The idea is that while the sun is shining, thermoacoustics is converting liquid water to solid water, storing "coolth" as the latent heat of a phase transition. At night, the ice keeps the medicine cold and when the sun rises again, you generate more ice. It has zero moving parts.

PARTICIPANT: How big is it?

DR. GARRETT: Overall size? Good question. It is about 60 cm. This has not been built. We have parts of it built but this is not yet running. I will give you some of the specs on the lower portion of this slide.

The reason I bring this up is that it has a prime mover, a motor, and a refrigerator, so it is time we can now look at what the laws of thermodynamics tell us about the overall performance of the system. There are some specs over here, but basically, we like to use these energy-flow diagrams. It makes things quite transparent, quite simple.

We have heat coming from the sun. We can measure the heat flux through this 3-foot diameter Fresnel lens. We have about 350 W on a good August day in Pennsylvania, both of them. (That shot was for Dick Stern's benefit, because I am always bitching about him moving me from California out to Pennsylvania.)

The heat flows through that prime mover, generates 50 W of sound energy. By the First Law of Thermodynamics, we must dump 300 W (the difference between 350 and 50) to this water-filled heat exchanger.

That acoustic work goes into this engine and it sucks 60 W of heat out of water to create ice. The 60 plus the 50 adds up to 110, which goes to this intermediate exchanger, again. The reason I picked this is -- notice -- this looks like it is a violation: I am getting 60 W of useful cooling for 50 W of electrical power. Is this a fraud?

[Transparency 25]

I want now to address the efficiency issue. These diagrams are from Matt Poesse's masters thesis. In a prime mover, we take some amount of heat, Q_{hot} , into the engine. We exhaust the work and we exhaust heat. The Q_{cold} goes to a cold reservoir, the Q_{hot} comes from a hot reservoir, we take out work.

The First Law tells us that the total energy that came out of the engine had better be the sum of the energy that left as heat and that left as work.

The Second Law of Thermodynamics tells us that the entropy that was created. The heat into the hot end divided by the temperature of the hot end, had better be bigger than the entropy absorbed by this reservoir, $Q_{\text{cold}}/T_{\text{cold}}$, because the entropy of the entire system must, at best, remain the same.

We define an efficiency in the sensible way. Since this is a motor, it is consuming thermal energy to give us mechanical work. The work is what we want. Q_{hot} is what we paid for getting that work out of the engine. So we define the efficiency as the work divided by Q_{hot} . That ratio,

called the efficiency, must be less than or equal to $(T_{\text{hot}} - T_{\text{cold}})/T_{\text{hot}}$. You can show that from the Second Law and the First Law combined.

The efficiency is obviously less than one, because the Third Law tells us that this that T_{cold} never reaches absolute zero. Nobody is upset with that. Nobody was upset when I claimed we get 50 W of work from 350 W of heat in the prime mover section of the solar ice machine.

Reh-lin's device is supposed to produce 50 W of work from 350 W of heat, which corresponds to an efficiency of 14 percent; 14 percent of the heat is converted to work, Carnot tells us that if the temperature of the hot end of the stack is 850° K, and the temperature of the water-filled heat exchanger is 300° K, a perfect Second Law engine (66) could give us 65% of the heat energy as work -- not 100%, but 65%.

Reh-lin claims that he will get 22% of Carnot efficiency. I suspect that he will not. He should get at least 10-15% of Carnot, but I am just showing you these numbers without all these little details that are going to drive down the efficiency of a real device.

[Transparency 26]

In the refrigerator portion we can do exactly the same thing, except here we are taking the work in, we are sucking heat out of a low-temperature reservoir, and the sum of those energies are getting dumped at some higher temperature. That is what a refrigerator or a heat pump is designed to do.

Here we do not define an efficiency, but a coefficient-of-performance, which is the same type of performance metric. We do not like to call it an efficiency because it is frequently greater than one. We want to suck Q_{cold} out of a cold reservoir. This W is the amount of work we must do to move that heat up to a higher temperature, T_{hot} . This ratio (67) must be less than or equal to $T_{\text{cold}}/(T_{\text{hot}} - T_{\text{cold}})$. T_{cold} will be about 270° K and $T_{\text{hot}} - T_{\text{cold}}$ is about 30° K. Therefore, Carnot says that we could take 9 times as much heat out of the cold reservoir as we are providing mechanical work to the engine. The coefficient-of-performance is bigger than one. It is generally bigger than one for modest temperature spans.

In your home refrigerator, the Carnot efficiency should be somewhere between 8 and 16, based on whether you are in Biloxi or whether you are in Monterey. Tom, this started with you. You started reporting this stuff on coefficient of performance relative to Carnot. Do other people do that, in general? I do not see it.

DR. HOFER: You remember that Colorado conference and they are talking in EERs, you know, which is not even in consistent units.

PARTICIPANT: What is EER?

DR. HOFER: It is the ratio of BTUs of heat pumping to watts of electrical power consumption.

DR. GARRETT: We should be thankful for the fact that thermoacoustics was started by physicists and, thanks to Tom's foresight, we have a measure that we call COPR, the coefficient-of-performance relative to Carnot. You take the COP, you divide it by what Carnot says you could get, and then you get something that is always less than one and we call it COPR.

If we look at Reh-lin's refrigerator, his solar-powered refrigerator, he expects to suck 60 W of heat out of the cold end with 50 W of work input, so his COP is 1.2., Over his temperature span (it is a larger temperature span) for the refrigerator portion alone, Carnot says the COP could be as large as 5.3. So, Reh-lin expects to see about 23% of Carnot in that device. I claim

he will see somewhere between 14% and 20% by the time it is actually working. Tom will back me up.

[Transparency 27]

This is a device that was used in Matt's masters thesis. It is known as "Frankenfridge." For those of you who do not know the work of Mary Shelley, I will provide a quote from the evil Dr. Frank: "I believe I can make a refrigerator out of parts of others." We have joined parts from two different thermoacoustic refrigerators. Due to a rare bit of foresight, this is a resonator that Tom and I built for a Space Shuttle experiment has the same bolt circle as the high-power drivers built for the SETAC experiment.

We joined the Space ThermoAcoustic Refrigerator resonator to a much more powerful loudspeaker so that we could explore this nonlinear regime that people were asking about earlier in this lecture. In-between the loudspeaker and the hot heat exchanger here, Matt placed a flange that had water running through it so that we could measure the exhaust heat very accurately. In fact, Matt could make measurements that were accurate to about 65 mW using this system; it was really very slick.

We can measure the acoustic power, W , very accurately since we know the acoustic power, the area of the piston, and the piston velocity from an accelerometer mounted directly on the piston. We know Q_{hot} from Matt's flange, so we can infer the heat load on the cold end without actually measuring the heat load on the cold end directly. Since we can apply a heat load with a Kapton electrical resistance heater at the cold end -- if we ever see that the cold end heat is less than the electrically applied heat, we know we have a problem. It works very nicely and you can get very good quality data very rapidly.

[Transparency 28]

I wanted to show you the measured performance of this 'fridge, because up until now, I have been fast and loose with the detailed theory. The linear theory has been put into a computer program that is a gem, called DELTAE, the Design Environment for Low-Amplitude ThermoAcoustic Engines. It was written by Greg Swift and Bill Ward at Los Alamos. You can get it for free and you can get the manual for free. It is available from the ThermoAcoustic Home Page maintained by Bill Ward at Los Alamos National Laboratory <www.rott.esa.lanl.gov>. That site also has links to our Penn State thermoacoustics web site as well as others such as the one at Purdue University.

You do need to be aware that DELTAE is a Ferarri and you do not want to learn to drive in this monster. It is very, very powerful. It solves the detailed thermoacoustic equations in each segment of the 'fridge. It is modular. Every DELTAE model starts with a BEGIN segment and ends with a segment called a HARDEND or a SOFTEND; a boundary condition that is "perfect."

In-between you then model the various segments of a complex device. You have an endcap, which is your piston. Then you have, if you look at Matt's fridge, something like this that has surface area and shape and then you have this and you have this, different parts. Those are a bunch of what we call ISODUCTS, isothermal ducts to model the bellows, the flange, and the tubes.

There is a heat exchanger, HXFRST. This is the first heat exchanger; it is segment 6, HXFRST, the hot heat exchanger. There is a stack, STAKSLAB, a stack of plates in segment 7. The cold heat exchanger is No. 8. Then we go from a neck to a cone to a neck, so there is

ISOCONE, INSDUCT. These are insulated segments that allow the heat that is produced in the neck to be dumped on the cold exchanger, which is the way it is in reality. Then there is a COMPLIANEC, which is this bulb. And then the final obligatory boundary condition, in this case, a HARDEND.

DELTAE integrates the pressure, the velocity, and the temperature along the model from BEGIN to the end. The pressure is a complex quantity, the velocity is a complex quantity, and the temperature is a real quantity, so there are five variables, five unknowns, there are five equations. This is a one-dimensional model. All of the action that is transverse to the acoustic axis is integrated algebraically in advance, so it is a very, very fast program.

It has a solver that is quite good. You give it a number of variables that you would like to hit, called "targets," say the heat pumping power or the temperature. You give it an equal number of "guesses," which are the variables that the program is allowed to twiddle to satisfy the boundary condition at the end while attempting to reach the specified "targets" as closely as possible.

DELTAE has a lot of useful "accessories." The model I showed you did not include the electrodynamic transducer, but it could. It has built into it the gas thermophysical properties. Remember, I used the short stack approximation; I said the temperature is uniform. I did not change the thermal conductivity, I did not change the viscosity, I did not change the heat capacities or anything else as the temperature changed along the stack. That is all done accurately, all that thermophysical parameters of the gas and solid as a function of temperature and pressure is in the computer model.

It was designed by users. This was designed as a product. It was designed as a tool for internal use at Los Alamos. It was started, actually, by Albert (Migliori), who was doing a prime mover design that used sodium-potassium eutectic as the "working fluid." It pays to actually do a few calculations before you decided to use something as a working fluid that turns water into fire -- lovely stuff that NaK.

DELTAE was developed by people who needed it; the documentation is excellent. The manual for DeltaE is the closest thing we have to a compendium of detailed theory on the behavior of working fluids, ducts, resonator dissipation, stacks, heat exchangers, drivers. I do not want to dwell on that; I want to dwell on the data just to show you that the theory is, in some areas, well in hand.

[Transparency 29]

These plots are taken from Matt Poesse's thesis with Frankenfridge, that device that you just saw in Slide 27. For those of you who do not know the book *Frankenstein*, it was the equivalent of *Jurassic Park* set back in a time when people were worried about surgeons and not about genetic engineers; when surgery was new.

If you look up here to pressure ratios of about $p_1/p_m = 3\%$, there are several things on here (and I am going to show you another graph but this one is kind of nice). There is a DELTAE model, which are the squares. There are the measurements, where you cannot see because the agreement is so close, at least up to pressure ratios of about 3%. The crosses are the DELTAE model with rectangular pores instead of just lateral plates.

I am going to blow up the portion that is below pressure ratios of 3% and let's concentrate on the lower graph. Do you know where this point is, Matt?

MR. POESE: It is not there.

DR. GARRETT: It is not there, okay, so there is a modeled point that is at a lower pressure ratio than he measured. You basically can see a difference between theory and experiment only at very, very low values of pressure ratio, where it is harder to measure the heat flows. Otherwise, the experimental data is covered completely by the DELTAE results out to 3 percent, and there are no adjustable parameters.

Another nice feature of DELTAE is a recent segment that they added that will take a block of data that you measured in your laboratory, and run DELTAE automatically for those specific operating conditions. What Matt has done is he has taken the conditions under which each data point was taken, or any of these other points were taken, ran those conditions into DELTAE and asked DELTAE to provide the theoretical cooling power he should have obtained for that pressure ratio.

DR. MIGLIORI: Is it a short stack?

DR. GARRETT: No, it was a damned long stack. It is a very, very bad 'fridge from the standpoint of efficiency. It is just a very good 'fridge from the standpoint of doing basic research. I do not know, what was your COP, 6 percent or something? Terrible.

It has a stack we had built in Monterey that Tom and I had done for getting to very low temperatures. The density of data here is enormous, as you can see. He took an awful lot of data and he took them very quickly, because he kept the temperatures very close to ambient.

In the upper graph, you can also see that above 3% pressure ratio, there is a systematic deviation of the data from the linear theory. The squares and crosses are generated by the DELTAE model results in the linear limit. The diamonds are the measurements. Here the data are on top of the theory and you can actually see that, it is not obscured. You can see it peeling off here above $p_1/p_m = 3\%$.

Remember, this portion here spans over two orders-of-magnitude in cooling power. Then up here, in the last factor-of-four increase in cooling power -- when you double the pressure you do not get four times the cooling power predicted by linear theory when the pressure ratio is doubled -- but you lose only about 15% from what you get in the linear model.

As I said, usually when you are designing a real practical device, weight, size, and power density are important. You may be willing, in many instances, to piss away 15% in efficiency to get four times the cooling power density. These were very encouraging measurements that we have submitted to Hank (Bass) and he will, I am sure, publish without refereeing, based on their importance.

[Transparency 30]

How does one design a 'fridge if you want to make one that would keep the commercial guys from doubling over with laughter? The useful heat-pumping power (Q_c), which I call $Q_{c, \text{useful}}$, is less than $Q_{c, \text{total}}$. You cannot sell all the cooling power. Some of it goes to cooling the dissipation on the cold end of your 'fridge, but they are roughly equal.

The net cooling power is the total energy you move, the enthalpy flux, minus the work you did to move it. From scaling arguments, it is fairly easy to show that the cooling power per unit area of stack is proportional to the pressure ratio squared, proportional to the mean pressure of the working fluid, and is also proportional to the sound speed in the working fluid.

If you are really sloppy, what you do is you write this down the product of those terms and then you divide that by the FOD (70). That is what I call the "figure-of-merit." For most devices that we have looked at, prime movers and refrigerators, FOD is between 40 and 120. It depends on how long your stack is, it depends on where your stack is located. It depends on a whole bunch of things. But basically the power density depends on your pressure, your sound speed, and your Mach number.

When you are faced with a design problem, you have to make certain choices. What is your highest acceptable mean gas pressure? The 'fridge that Tom and I put on the Space Shuttle in January, 1992, worked at about 11 atm. I was then gutsy enough to go up to 20 atm on a 'fridge that went out on a Navy destroyer. I will describe it if I have enough time near the end. The ones that we were building now, the large air-conditioner, is up at a mean pressure of 30 atm. We keep increasing the pressure because it increases the cooling power density (70).

DR. MIGLIORI: Does DELTAE include streaming and kinetic energy effects?

DR. GARRETT: No. Bill and Greg are throwing in more components like that. They have added certain turbulence things from D.C. flow that appear to be mimicked by A.C. flow. You now have a choice with an INDUCT or an INSCONE to provide a surface roughness. If you set the surface roughness to zero, DELTAE gives the linear thermoacoustic result.

If you set that roughness to be something else, it will go through all this sort of Moody diagram, turbulence correlation stuff that is out in there in the mechanical engineering literature.

That is exactly the right question to ask next, Albert, that is at how high a pressure ratio do you think your equations are going to work? That is another decision you are going to make. Are you going to design for a pressure ratio of 3%, 4%, 4.5%, 5%, 6%, more than that? It is a decision you have to make.

DR. MIGLIORI: It is not satisfying to have that 15% floating around. You know almost everything you need to know. I remember working the whole streaming problem but I never put it into the code.

DR. GARRETT: These are important issues, definitely important issues, but it is not stopping us from making locomotives.

You can neck down the resonator -- a Helmholtz resonator is a lot shorter than a standing wave resonator for a given frequency. To decrease the size for a given frequency, you increase the velocities and you increase this risk that Albert is worried about.

What trades are you willing to make for that efficiency? If you go to a lower Prandtl Number, your efficiency goes up. That is wonderful, but your power density goes down, because lower Prandtl Numbers correspond to lower sound speeds. So this is this "clanking hissing" business that Prof. Wheatley mentioned. Multiple stacks or single stacks? Multiple heat exchangers or single heat exchangers?

[Transparency 31]

Much of the tradeoff is in stack location. If you put the stack very far from the rigid end, you have a lot of velocity, so you get a lot of cooling power, but as you bring the stack in, your coefficient of performance increases because the a shorter stack can span the same temperature difference. Less stack, less loss. So, where you locate the stack in the standing wave is important. The guys at Purdue went for the maximum efficiency and they got a very big 'fridge. There is some tradeoff.

[Transparency 32]

The performance for a given stack in a given location depends on the temperature span. Remember the $\Gamma-1$. When Γ was zero, you achieve the maximum heat-pumping power. That occurs where the temperature span is zero, so you pump the most amount of heat, but at a very low temperature span Carnot is enormous. Even though the heat pumping is the highest, your coefficient-of-performance, divided by Carnot, is the lowest.

Down at the high temperature span end of the performance curve, the temperature span is very large but the heat pumping has gone away. Your coefficient-of-performance is zero and so is your coefficient-of-performance relative to Carnot. Somewhere between these two zeros is a maximum.

At first cut, that is the way you would design a fridge. You try to figure out what the necessary temperature span is, move your stack around so that you have a pretty good performance over your intended operating temperature range. In this case, you would like the coefficient-of-performance relative to Carnot that is about 30%. I would also like a call from the King of Sweden. I am sure that neither will happen.

The reason is that this a stack calculation and it does not include the resonator and heat exchanger losses, it does not include the turbulence, but we do get 30% of Carnot in our best designs.

[Transparency 33 and 34]

Let me show you a 'fridge that was designed to do the kinds of things that real 'fridges are supposed to do. This is the 'fridge that went out on the USS Deyo, which is a destroyer, but it looks nothing like that diagram from the moving parcel picture in Slide 15. It is thin in the middle, it is bent in half, and it has two loudspeakers. I provided this drawing because it has numbers and the numbers have a legend.

It is a half-wavelength device, so the gas is sloshing the U-shaped section. There are two stacks. Because there are two rigid ends, you might as well use both pressure antinodes; it can make the resonator diameter smaller by $\sqrt{2}$. It has two 11 cm diameter stacks, just a roll-up stack, four heat exchangers with fluid-filled backbones, and two custom-made loudspeakers. If you have a half-wavelength standing wave mode in the resonator, the individual drivers are moving 180° out-of-phase.

Of course, you could do it with a single driver, because a double-acting piston has the right phasing to do it, also. But this 'fridge was originally designed for NASA, and they always want the redundancy.

Under the highest loading conditions, the SETAC 'fridge was generating 216 W of acoustic power and moving 419 W of thermal power from the radar units onboard the USS Deyo. The loudspeakers had about 52% electro-acoustic conversion efficiency.

The heat exchangers were underdesigned since this was supposed to be a small space-based refrigerator for life science experiments. It was supposed to move only 100 W and here it is moving 400 W through the cold heat exchangers and exhausting 600 W out of the hot heat exchangers. You can see there is quite a bit of temperature difference across the stack in order to suck that much heat through these underdesigned heat exchangers.

The overall COP, based on the reservoir temperatures, is only 8%; not very impressive. The engine itself, which was designed to work at a much higher ΔT , is performing at a reasonably respectable 21% of Carnot.

DR. MIGLIORI: What was the predicted?

DR. GARRETT: This is very close to predicted. It was predicted to peak at 30% of Carnot, and we have now we can run it with methanol.

The coldest we could get with water on the cold side, because the radars are cooled by water, they are not cooled with methanol, and so if you cool too much, you freeze your heat-transport fluid. This is the coldest we could get sucking 300 W out and putting 166 W of acoustic power in.

You can see that even a small reduction in the load on those heat exchangers kicks the overall COP relative to Carnot outside the engine up by a factor of two. It went from 8% to 17% of the Carnot. In addition, the engine COP, of course, went up, also, to 26% of Carnot.

[Transparency 35]

In my remaining four minutes, I am going to say something about the design. Remember, the cross-sectional area of this resonator was not uniform. Since Mark Hamilton talked about harmonic suppression, I should point out that we have the same problem. The sound levels in there are 190 dB. 160 dB will light your hair. These are pretty high sound levels! You are not worried about waveform shocking, you are well past where shocking would begin.

What we do is we change the cross-section of the resonator like the MacroSonix boys. It is easy to see. An alternative way of looking at that is if you have a uniform resonator, as shown under (71) the frequencies of the overtones are integer multiples of the sound speed over twice the resonator length.

If you put an obstacle, a solid particle, at one end, you basically shorten up the resonator. If you think of this solid block as wax and melt it down at one end, you have made a shorter resonator. This increases the frequency of the fundamental mode.

If you put the obstacle in the middle of the resonator, the kinetic energy of the gas has to increase as it speeds up to get around the blockage. This drives the frequency of the fundamental mode down. If you leave the block in the middle and look at the effect on the frequency of the second mode, the block is at an anti-node, again, and it drives the frequency of the second mode up. By changing the cross-sectional area, Mark told you it introduced dispersion through the horn equation, I am telling you in a much more simplistic way that these perturbations in the geometry detune the overtones from the fundamental. Resonator shape is important, because it can suppress the nonlinear generation of harmonics that would limit the achievable sound levels.

[Transparency 36]

You can make very simple models for what area the pistons should be. You know the work generated to pump the heat is the pressure times the volume velocity divided by two (72). These are peak quantities. Pressure times the volume velocity is equal to the pressure times the piston velocity times the piston area. Pressure times area is force, so we know force is $B\dot{V}$, for an electrodynamic driver.

The power dissipation in the electrical resistance of the voice coil is $R_{dc}i^2/2$, again, using peak quantities. It is proportional to the area squared, because as you make the area larger for a

given acoustic pressure amplitude, you have to use more force and more force requires more current.

On the other hand, the mechanical dissipation is proportional to $R_m v^2$. If you use more stroke and less area, you generate the same amount of acoustic power, but more stroke means more velocity, and that drives the mechanical dissipation up. That loss goes as the reciprocal of the piston area squared. If these are your only two loss mechanisms, there is an optimum area give in (75) where that area corresponds to equal loss by both mechanisms.

That analysis is a bit simplistic. Robert Keolian can do it right. You actually have three coupled-oscillator systems. The inductance to the voice coil, its resistance, and maybe a power factor correction capacitor form an electrical resonator. That electrical resonator is coupled to the mechanical resonance of the driver which is determined by the suspension stiffness and the moving driver mass. That mechanical resonator is coupled to the acoustic resonance. Robert has modeled the coupling between all three of them. You have got to do it correctly if you want to squeeze all the efficiency out of the system, but the approach in Slide 36 gets you very close if you have a reasonably efficient driver.

[Transparency 37]

I prove here that the dissipation in the voice coil depends only on the weight of copper that you are willing to use. You saw in the earlier loudspeaker, we had copper on the voice coil which was moving with the piston. If you increase that copper, you are going to have to use more force to accelerate and decelerate the additional mass, or you will require a stiffer suspension to resonate away the effects of the additional mass.

[Transparency 38]

The new refrigerator that we are working on inverts that. This linear motor was invented by John Corey of Clever Fellows. Here the magnets are in the moving frame and you can put as much copper as you like on the stationary pole pieces. The electromechanical conversion efficiency of this driver is nearly 90%.

It is very pretty. It is potentially very easy to manufacture. You have these pole pieces. You can slide the coils in, you do not have to wind them *in situ* -- that drives the price way down. The elasticity comes mostly the restoring force of the magnets. If you try to pull the magnet out of that, the iron is going to suck itself back in.

[Transparency 39]

What does it drive? Seven feet long, 10 kW of cooling capacity from 4 kW of acoustic power. It is the 700-lb gorilla, except it is 2,300 lb. Just to show you that this is not a "pipe dream," below the artist's rendering is the actual hardware. It is a clever design. There are large flanges that squeeze the stack, the heat exchanger, and something I am not going to tell you about in here, so you can make modification to the internal sub-systems without too much difficulty. It should be a very good research device.

[Transparency 40]

In this view, you can see the inside of the resonator and the driver suspended by a ceiling crane.

[Transparency 41]

Let me show you a close-up of some of the parts that have been built. I like this picture, it is called "3 heads." Tom Gabrielson, head number one back there, is the designer. He happens

to be the only guy in the thermoacoustics group at Penn State who has a professional engineer's license. He had actually seen Goodman graphs and ASME boiler codes, so he had the privilege of designing the pressure vessel. We pressure-tested it up to 900 psi -- he is still alive, so we know that he did a good job. These bolts are an inch and a quarter in diameter. The flange is enormous because of the large moment. The stacks and the heat exchangers squeeze into the space between the. Its useful internal (stack) diameter is about 19 inches.

With that, I want to point out that I have left much good work in thermoacoustics out of this lecture. I am glad to be able to make that statement; there is a lot of work going on in thermoacoustics all over the world. In Japan, there is a society of thermoacoustic heat transport. It has got a hundred people, half of them have Ph.D.'s. Greg Swift and his large group at Los Alamos is doing wonderful work using heat-driven pulse-tube refrigerators.

Logan is standing up there, so I am going to put this up and call it quits. I would just like to close by stating that the industrial revolution owes its success to the fact that the computer had not been invented yet. If it had, we would still be modeling and simulating the cotton gin, the telegraph, the steam engine, and the railroad. Thank you.

Date sent: Tue, 07 Dec 1999 09:11:23 -0600
To: lsavage
From: "Henry E. Bass" <pabass@sunset.backbone.olemiss.edu>
Subject: Fwd: EXPENDITURE RATE & FY 2000 PARI FUNDS heb

>From: "Hargrove, Logan" <HARGROL@ONR.NAVY.MIL>
>To: "'BASS'" <pabass@olemiss.edu>
>Subject: EXPENDITURE RATE & FY 2000 PARI FUNDS heb
>Date: Tue, 7 Dec 1999 09:57:42 -0500
>X-Mailer: Internet Mail Service (5.5.2650.10)
>
>HANK,
>
>\$1M has recently appeared in my "cash box" for PARI,
>without my having to "negotiate" with Fred! That does
>not mean he wants me to send you money.
>
>The CAMIS Report I ran this morning shows the last billing
>to be Voucher No. 11, received 18 OCT, dated 11 OCT,
>covering the period month of SEP, and in the amount \$9,076.72. This
>brings the billing against N000149810033 FY99 funds (\$1M) to
>\$510,353.01 which is 51%. End-of-December Benchmark for prior-year
>funds is 80% EXPENDED. All you have billed has been paid and posted
>in STARS as of last Disbursement date of 28 OCT.
>
>It appears that U of MS is a bit behind on billing.
>I hope we can expect a billing for the months of
>OCT and NOV soon, followed by a timely billing for DEC.
>
>I will have to make fewer arguments up the approval chain to PR the
>\$1M of FY 2000 funds if your FY99 Expenditure looks better.
>
>I recall that you said you did not expect to need additional
>funds any earlier than about FEB 00.
>Does this projection still hold?
>
>Maybe I sound like Scrooge to bring this up when the Holiday Season
>is upon us, but I remind you that in the end, Scrooge laughed.
>
>LOGAN
>
>DR LOGAN E HARGROVE
>ONR 331 ROOM 503-13
>OFFICE OF NAVAL RESEARCH
>800 NORTH QUINCY STREET
>ARLINGTON VA 22217-5660
>
>Voice/Voicemail (703) 696-4221
>FAX (703) 696-6887
>

NONLINEAR ACOUSTICS

Mark F. Hamilton
Department of Mechanical Engineering
University of Texas

DR. HAMILTON: I was looking at the program for the upcoming meeting of the Acoustical Society in Seattle next week and I notice there are six different sessions on nonlinear acoustics there: two will be devoted to fluids, two to solids, and two to biomedical applications.

The field of nonlinear acoustics is quite broad. Some people would claim that the term "nonlinear acoustics" is something of an oxymoron, that when we talk about the acoustical approximation we often refer to the linearization of exactly these equations that contribute to these nonlinear effects. The word "nonlinear" is, of course, the mathematical description of the equations that are used to describe the sound field.

It is reasonable, when you do the linearization of these equations, you take, for example, in your equations of state or equations of motion, your field variables, perturb them about an ambient state, and then linearize with respect to the perturbation of these quantities.

In fact, for the wide variety of cases in acoustics, this is a completely valid approximation. Even if you look numerically at these quadratic and higher order terms that are neglected, you will find that numerically they are indeed very, very small terms.

For example, if you consider a sound wave in air at about 154 dB, that would have an acoustic Mach number of 0.01, the acoustic Mach number being the particle velocity in the traveling wave divided by the sound speed. This 1% indicates the relative order of these nonlinear terms, these quadratic terms in these equations, relative to the linear terms.

As a first cut, when you look at these equations, you might automatically, in many analyses, dismiss such terms. However, this same sound wave at 150 dB, if you drive that at 1 kHz and transmit it as a plane wave, you will find that this wave will form a shock front on the order of 5-6 m, so in fact, the process can be substantially nonlinear, even though these small terms that we often disregard can be very small.

Why does this happen? The reason it happens is because the way the nonlinearity enters the equation is in a cumulative sense. It may be small locally, on the scale of a wavelength, but as this wave propagates over many wavelengths, we have an accumulation of these effects and these little effects added can give you substantial changes in the waveform.

[Transparency 1]

Here is an example of such a case. The main effect, if you look at the nonlinear propagation of a sound wave, of a compressional wave in a liquid or a solid, you will find that, in general, the wave speed along the waveform changes from point to point. This is, in general, a very small correction. This additional correction to the sound speed here is on the order of, say, 1%.

However, as the wave propagates, this peak tends to catch up with this zero crossing, this trough tends to fall a little bit behind, and ultimately, on the order of many, many wavelengths, one gets distortion of a waveform that can eventually end up with a shock front. At this point you need to take into account dissipative mechanisms to describe the structure of that sound

wave, but this gives you an idea of the global types of effects you experience in nonlinear acoustics.

Because the equations are nonlinear, in general, they are exceedingly difficult to solve, in particular to get exact analytic expressions of these equations. To look at any specific problem in detail often requires either tricks to be able to solve these nonlinear equations or extensive approximate methods, asymptotic perturbation methods, for example, to obtain solutions.

Rather than dive into any one of these areas specifically, the purpose of this talk is to give a survey of the types of basic phenomena and to introduce you to the working language and the physical concepts that are involved in the processes of nonlinear wave propagation.

The types of topics we will be considering initially will be waveform distortion and, as this wave distorts, there will be new frequencies generated in this wave. Shock waves will be produced and we will also look a little bit at some of the D.C. effects, namely, radiation pressures and acoustical streaming -- these are the time-averaged properties of sound fields that result from these nonlinear quantities.

[Transparency 2]

The study, theoretically, of nonlinear acoustics dates back, really, 2 centuries to the origins of the studies of linear acoustics itself. In fact, the first equations of acoustics, for the most part, that were derived were originally nonlinear. The theorists got stuck, around the early 1800s, about the point of Poisson here. He was the first to be able to bring the equations of motion for a perfect isentropic gas into a very succinct and elegant form that lends itself quite readily to physical interpretation. In fact, we use his equation still today, and I would like to review that, because the main features of nonlinearity come from this.

We will begin by considering a perfect fluid with no energy losses; therefore, we will not require an energy equation. Consequently, in our equation of state we can relate the material properties of a fluid by just 2 variables. We choose the pressure and the density of a fluid. We will also consider an adiabatic gas with a gamma, which could be 1 for an isothermal, so this is arbitrary.

These are the exact equations of mass conservation and momentum conservation for a 1-dimensional disturbance in a gas. It is by no mean obvious, looking at these equations, how to do this, but by some very clever substitutions one can find that an equivalent rendering of these 3 initial equations is the following very simple form. This equation is exact for perfect lossless gas for a sound wave that is propagating from the left to the right in one dimension.

Here U is the particle velocity. C_0 is what we refer to as the small signal sound speed. We make a distinction here between small signal or the phase speed of the wave in the linear approximation and the finite amplitude sound speed that will vary from point to point along the waveform, as I showed in the previous slide.

The other coefficient that will appear prominently in our discussion is labeled β and is referred to as the coefficient of nonlinearity. Notice that if β is 0, this term vanishes. Here we have the only nonlinear contribution to this equation and the result is linearized, $du/dt + C_0 dU/dx = 0$, and one will recognize from that the simple traveling wave solution of $f(x) - ct$ for a wave traveling from left to right, the d'Alembert linear solution. This equation resembles very closely what we are used to seeing in the linear approximation.

Another point to notice here is you can examine, at least in a cursory way, where the nonlinearity in these equations comes from, but back to the equation of state for a moment. Here we have a nonlinear relation between the pressure and the density, and γ , for example, is 1.4 for a diatomic gas, which is usually the relation that we use for air.

If you have a circumstance where the gas is isothermal, the temperature is constant from the ideal equation of state we have and the pressure is proportional to the density, in this case, γ would be equal to unity. Notice that the coefficient of nonlinearity, then, reduces to 1. We still have nonlinearity in these equations by virtue of the continuity in the momentum equations.

Here is our progressive wave equation describing a wave going from left to right along the x axis. An exact implicit solution of this, which Poisson found and bears his name, is given by this expression. Again, it resembles very closely the equation of linear acoustics. In fact, if β were 0, we have $t-x/c$, and that is our traveling wave.

However, we have, in this circumstance, a wave field that is a function of itself, so we have an implicit function that is described in the propagation of this wave. Although we have an exact solution, it may be difficult to implement. In this case it turns out that the equation can be implemented quite easily in a graphical manner.

That is to say, if we take a look at this equation and fix U and ask how fast that portion of a wave, a specific part of a wave that we tag on the particle velocity, propagates as a function of time, this propagation speed, $dx/dt = c_0 + \beta U$ -- you can see it by inspection here in the denominator. This is what gives rise to these cumulative nonlinear effects in the progressive wave distortion that I showed on the previous slide.

[Inaudible question]

DR. HAMILTON: Beta, actually, is a more fundamental thermodynamic quantity you can derive. Beta has been found to be negative, not in any gases and, to my knowledge, in liquids only near critical points, and in superfluids. In a gas it cannot happen and in liquids in normal states it cannot happen.

DR. GARRETT: In a gas, γ can go only between 1 and 5/3 --

DR. HAMILTON: In solids you can get negative coefficients of nonlinearity. You do not have exact solutions like this, but because of the constitutive relations of the solid -- for example, in the solid you can construct a situation where you can actually have a negative Poisson ratio, if you are clever about it; you push it and it squeezes in. There is an example I have heard used, origami constructions, where you can push in a kind of lattice and it will come in.

MR. HICKEY: What kind of material has a beta equal to 0?

DR. HAMILTON: In solids, I know, for compression waves, glasses have negative coefficients of nonlinearity. These will shock backward from what we are familiar with in liquids. For example, fused quartz is one example. Pyrex is another example. The compression wave will steepen in the opposite direction.

DR. GARRETT: Cars on a highway. Thermal waves in superfluid get you into 1.884 K as opposed to 0.

DR. HAMILTON: This may address somewhat your question of negative nonlinearity. In general, if we relax our restriction of a perfect gas, we have to appeal, then, to a more general formulation of the equation of state. Again, we will ignore losses here (we will consider the loss

terms later), so we consider the pressure as an explicit function of the density and perform a regular Taylor series expansion about its ambient state.

[Transparency 3]

One can then proceed to write such an expansion, introduce the perturbations of the pressure and the density in terms of the acoustical primed variables here, and rewrite the expression like this (this is a notation made popular by Bob Beyer), and we identify the linear quadratic and higher order terms, then, with coefficients as in this Taylor series expansion that are traditionally now called simply A, B, C, et cetera.

The relative nonlinearity in this equation of state for the fluid is given by the ratio of the quantity B/A . B/A , just taking a look at the basic definitions from the derivatives, is given by this quantity. C_0^2 we recognize as our isentropic small signal sound speed, the slope of the pressure density relation at the operating position of the wave at the ambient state in the fluid.

Again, here you can see in this ordering procedure where the smallness of these terms come into that I was referring to previously. For that sound wave of 154 dB, where the Mach number was 0.01, 1%, that is effectively the same as the amplitude of this quantity, ρ'/ρ_0 . You can see even in this example here in the equation of state expansion this term at 150 dB in air, let's say, is 2 orders of magnitude smaller than this. Nonetheless, it is an important quantity in a cumulative sense as the wave propagates.

What types of values do we have for this parameter? For most conventional liquids B/A ranges in the vicinity of 5 to 10 and for water it is 5.0, very close to water at room temperature. If you were to equate this expansion with the same expansion of the isentropic equation of state for a perfect gas, you would find that B/A corresponds to quantity $\gamma-1$, so if you were writing it in this notation, B/A for gas is 0.4.

Let's go back to a physical interpretation here. We take a look, again, at the propagation speed and we recall that from the Poisson solution it is given by the small signal sound speed C_0 , with the correction β times U , where U is the local value of the particle velocity at any given point on the waveform. Beta in this generalized case is $1+B/2A$. With $\gamma-1/2$ placed in here, we have our $\gamma+1/2$ back for the perfect gas.

Written this way, note the interpretation. If we have a fluid in which the constitutive relation is linear, P is linearly proportional to ρ , for example, an isothermal gas, then B/A is 0 in this case and beta is unity. In this case, the correction to the local small signal sound speed from point to point on the waveform is exactly the particle velocity U of the wave. We refer to this first term in the coefficient of nonlinearity as the convection term. The wave is convecting itself.

As that wave propagates, the pressure corresponds to a particle velocity that is in phase with that wave for a traveling wave to the right and that particle velocity then convects the wave with it; it actually rides along with its own flow. To a large extent -- we will see this later -- this convection term comes primarily from the mass conservation equation. Even though we refer to it as the convection term, it is not from the momentum equation.

This is the general formulation for how to characterize fluids and you will see much discussion about the parameter B/A in the literature and a lot of investigations have gone into how do you measure B/A to characterize the nonlinearity of various liquids and gases and, more recently, solids, which can be extraordinarily nonlinear.

I will skip the next slide, which is waveform distortion, which we have already seen, and say just a few more things about measurements, how do you get at this parameter B/A in a liquid, how do you get values for these.

[Transparency 4]

The basic definitions of these fundamental constants that we obtained from the straightforward Taylor expansion of the equation of state are derivatives of pressure with respect to density, so B/A , then, involves taking 2 derivatives of pressure with respect to density under isentropic conditions. Well, imagine this for a liquid.

What you are asked to do is vary the density of a relatively incompressible substance and then try to characterize the variations of the pressure with respect to those changes in density; not a practicable way of doing the measurement.

An alternative way, one of the ways that is used, is through thermodynamic identities, the ratio B/A can, instead -- in fact, you can see this immediately -- we recognize that the sound speed is $dp/d\rho$, so one of these derivatives inside here we can replace with C and we can write this as a derivative of sound speed with respect to density and then, using another relation for the sound speed, write it as the derivative of the sound speed with respect to pressure under isentropic conditions. This is much easier.

You are changing the pressure now but it must be done adiabatically and at the same time monitoring the change in the sound speed. If you are doing small signal experiments, we are propagating, for example, in an electrosonic beam through a liquid, which would presume you could perform that experiment rapidly enough that isentropic conditions are maintained. This can be difficult unless you make those changes relatively rapidly and smoothly.

To do that, you can actually use the sound wave itself. As the sound wave propagates through, it itself is a modulation of pressure. By cleverly recognizing the phase modulation of a wave through this relation, one can back out the values of B/A . Again, this is somewhat tricky.

C/A , if the quadratic term is small, the cubic term is very, very small; nevertheless, we have an expression, also, that we can write for C/A in terms of variation of the sound speed with pressure, depending also on B/A itself.

Just one comment about that. Despite the fact that I said these nonlinear effects are quadratic order, even though numerically small, and can be quite substantial for long distance propagation. The term C/A , the cubic term in the equation of state, has very little impact and is almost immeasurable in most circumstances for normal liquids, solids, and gases, so very little has been reported on the value of C/A and, in fact, its effects are very, very small.

Finally, and the more conventional way of measuring B/A , is by yet another thermodynamic transformation of these relations for B/A , in which we now vary the sound speed as a function of pressure, holding either temperature or pressure constant. This, in the literature, is referred to as the thermodynamic method and the bulk of the data that are out there is primarily based on this method.

The relative order of these terms, incidentally, for most liquids, is the first term dominates the second in this expression by a factor of about 3 or 4 to 1, usually. It is the variation of sound speed with pressure that is the dominant nonlinear contribution. You can read about this in this article by Coppens, et al.

[Transparency 5]

Some typical values of B/A for common liquids -- I show this only to give you an idea of the span of ranges it is for most liquids one will encounter. For water you see that the changes are relatively small and we usually use 5.0 for distilled water at room temperature for the coefficient of nonlinearity.

For other liquids this ratio varies on the order of between 5 and 10 and over here you see a compilation for the equivalent B/A for gases. Ordinarily we do not quote B/A for gases; we think of it more in terms of γ in the isentropic equation of state equation relation. This is a table from Bob Beyer's 1974 book on nonlinear acoustics.

Incidentally, I gave you a preprint -- I forgot to mention this -- in the preliminary material that was sent out and it is a short review article. Beside just being a simple overview, perhaps its main utility is it has got lots of references, so it is a good bibliography of a lot of the classical work for these references I am citing here.

DR. MAYNARD: What is B/Z for liquid helium?

DR. GARRETT: We did not think about it in terms of B/A . It is more nonlinear than water.

DR. HAMILTON: We do very little low-temperature work in Texas. (Laughter)

[Transparency 6]

A hot area currently -- in fact, as I mentioned, there are 2 sessions on high-intensity ultrasound in biomedical applications at the Seattle meeting -- is biomedicine, so a considerable amount of work has gone into characterizing tissue samples by B/A , and this gives you an idea of the range of values there. Again, it goes on the order of 5 to 10, there is nothing particularly unusual about it.

You see here, from this column (this was taken from a paper by Bjørnø) -- the 2 methods used here are in this column. The finite amplitude method refers to the first of the relations for B/A on the previous slide. The thermodynamic method refers to the second of these. In the case of tissues here, the finite amplitude effect is used more predominantly.

A comment on one of the initial reasons for characterizing some of these tissues: A lot of work done, for example, at Yale University and up at the Mayo Clinic, had to do with the development of mixture laws.

DR. KEOLIAN: One of these things on this table is solids and your development was for fluids.

DR. HAMILTON: Solid-like, yes.

DR. KEOLIAN: Yes, and your pressure is a fluid concept. Are you approximating these solid --

DR. HAMILTON: Yes, there is an assumption, they are assumed to be liquid-like, sufficiently low shear. I will give you B/A later for solids, though. But you are right, it is a completely different formalism.

For most intents and purposes the characterization of body tissue as fluid-like is seemingly reasonable, although there is work now with low-frequency shear waves in tissue for imaging, generated, for example, by lasers and other methods and the wave speed is very, very low in this case, tens of meters per second or something like that, very low shear-wave velocities.

One point I will mention here is that in the tissue within the body, what you will notice, going from tissue to tissue, is the value of B/A in terms of its relative variation could be of order

1. You can go from 5 to 7, 5 to 8, going through layers of tissue, whereas if you are looking at the variation of linear effects in these layers, say, for example, the sound speed or the density change that you would use for imaging based on impedance mismatches for reflection, these are much smaller relative changes, on the order of a few percent, whereas you have a few percent change in linear effect, you have a many-order-1 change, often, in the nonlinear effects.

This got a lot of people's attention and hence a lot of these tables, because they thought, aha, we have a new imaging modality we can use here for diagnostics. Of course, the problem is you are looking for a big change in a small quantity, so there is always that compromise. The jury is still out, to some extent, on that, although a number of the papers in Seattle next week have to do with nonlinear effects used for imaging purposes.

[Transparency 7]

Let me give you kind of a cute example, one that Bob Keolian will be doing, I believe, for one of his demonstrations. This is an effect called the suppression of sound by sound. What is clever about it is it really brings to life what happens with this waveform distortion. You can see, just phenomenologically, what is happening and become a believer of really how these waves propagate.

Suppose we have a source that radiates at two frequencies and let us suppose that it is a very low-frequency, high-amplitude wave and the other frequency is a relatively low high-frequency wave. If we look at the initial spectral lines, Ω is the low-frequency, so-called pump wave. That one will be the one to carry the main nonlinear contribution. The high-frequency wave we refer to here as ω . That will be the low-frequency wave, sometimes, in applications like this, called the probe wave.

The idea here is to investigate how the high-amplitude, low-frequency pump wave modulates and distorts the high-frequency probe wave. Imagine what is going to happen here. The superposition of these 2 waves at the source looks like this. Here is my low-frequency pump, my high-frequency carrier riding on top of it.

Now let's focus on what is happening to the high-frequency carrier. You can think of these as almost like markers tagging different points on the waveform of the low-frequency pump. As this wave propagates, the waveform, because of the nonlinear effects I mentioned, will start steepening forward. That means there will be a compression here in the location of a shock where these high-frequency cycles will be shoved together, and along the more expanded section here they are going to be stretched apart a little bit.

If we now filter the high-frequency wave out of the composite signal, we will find a waveform that looks like this, cycles scrunched together here, spread out over here. To a first approximation -- what this is, of course, is frequency modulation of this high-frequency wave -- what you find is that the phase modulation of this wave is sinusoidal to dominant order.

It appears, mathematically, like sinusoidal FM modulation of one wave by another and for that we have an exact solution. There is a solution of that modulation that is given in terms of Bessel functions and that would describe all of these side bands. In other words, the FM modulation of this frequency ω by this low-frequency Ω will give rise to the side-band spectrum, so harmonics and side bands will be generated because of this distortion.

What becomes interesting, then, or novel, is if you then look at the wave at the probe frequency ω . It is described by a Bessel function of order 0 and the argument of that Bessel

function depends on 2 parameters of the primary wave field. It depends on its frequency but it does not depend on its amplitude; it depends on the amplitude of the pump wave, the low-frequency wave.

Bessel functions go through 0's, so if you take a look at this argument, you could say, well, I am going to fix x at some point in space, and ω is fixed, because I know what the frequency of my probe is, and I am going to vary the amplitude of my pump wave. As this happens, you are going to slide down this curve, passing this 0 at 2.4, and this side band will go down and disappear at this point.

Sometimes this phenomenon is called absorption of sound by sound, because you are using this sound wave here to take energy out of this spectral line, but there is no absorption in this problem, we have not invoked any dissipation mechanisms here, so we prefer to use the term "suppression of sound by sound." All you are doing is taking the energy out of this spectral line and spreading it into these side bands.

[Transparency 8]

It really works, it is easy to show. These are some waveforms that one of the students at Texas took, a former summer school graduate (of the first school, actually). You cannot even see the high-frequency wave and the composite wave here, so you have to go in and extract the high-frequency wave and, sure enough, here is that phase modulation. Here it is where the wave is shocking up and you can see the signals are really rammed together here.

[Transparency 9]

Take a look at the spectra -- this viewgraph did not come out well. I took this from his thesis and, unfortunately, it was color-coded and did not come out well. These are the spectra that resulted. In the upper left is close to the source pump wave at 950 Hz and sound pressure level of 140 dB, so, as I mentioned, at about a 1-kHz wave at about 150 dB we would expect a shock at about 5 m to give you a feeling for the nonlinearity of the system here.

As the pump propagates, it generates its own harmonics, but what we are interested in is the high-frequency signal here at 67 Hz and, as you propagate along (it is hard to see in here) the signal goes down and, as that goes down, it is generating all these side bands about it according to this FM modulation scheme.

[Transparency 10]

In that review paper I handed you, this is one experiment that is reported that was done in China. The suppression here, you can get it down to about 40-50 dB or so. This is not exactly what you would call a noise-controlled tool, because all you are doing is spraying energy over your frequency spectrum, but it is a very nice illustration of this distortion phenomenon.

MR. HODGSON: You have a certain amount of acoustical energy at the low-frequency pump wave. You have a certain amount of energy at the high frequency. What is the balance of energy during this process? Are you taking energy from the pump wave up into the neighborhood of the high-frequency wave?

DR. HAMILTON: It is conserved, first of all, prior to shock formation in the absence of linear effects, so the sum of the squares of all of these spectral lines is a constant. That gets to be a more delicate issue -- we will come back to it -- because as you start pumping this energy up into higher and higher frequencies, what seemed to be a relatively lossless fluid down here is a

relatively lossy fluid up here, because your absorption has gone up as, on the free field, say, ω^2 , so eventually it will start stripping off at this end.

Fluid ultimately looks like a high low-pass filter and strips away your higher harmonics, so it is shifting up and then being sucked out by viscous and thermal loss mechanisms.

MR. HODGSON: So how long a propagation distance can you go before you stop shifting energy from the low frequencies up to the high frequencies?

DR. HAMILTON: We have a number that describes that, it is a dimensionless quantity. It gives you the ratio of the amplitude of the wave to the viscosity of the fluid, so it depends. If it is a very strong wave, very far; if it is a very weak wave, not very far.

MR. HODGSON: But it has to do with the ratio of the amplitude of the pump wave and the viscosity?

DR. HAMILTON: That is correct. Because these are the competing mechanisms. Nonlinearity is trying to steepen the wave forward, dissipation is trying to hold it back. It is that balance of which one is going to win and where.

DR. GARRETT: Mark, the Goldberg number is nondimensional, right? and it does not depend on the amplitude, and the reason is you are taking the ratio --

DR. HAMILTON: The Goldberg does not.

DR. GARRETT: The Goldberg number does not depend on the amplitude?

DR. HAMILTON: Yes, it does. It is the amplitude of the wave over the absorption coefficient.

DR. GARRETT: Oh, because the shock-formation distance depends on it, yes, you are right. Sorry.

[Transparency 11]

DR. HAMILTON: This is a little out of order. In fact, if you go to the very end of your slides, I am jumping to sound beams now because of some discussions I had about a recent application of nonlinear acoustics called parametric array in air, and I felt this was a good place to move it up to and say a few words about it, because it fits rather well into the harmonic generation I just referred to.

At some point you have got to solve these equations. Pictures are nice, but at some point you need numbers to get a quantitative description of what is going on here. Let's, as a model equation just for discussion, use the Poisson equation that I gave you earlier. Remember, the linear terms $dU/dt + C_0 dU/dx$ for a traveling wave going from left to right, and there was a single nonlinear term that was quadratic in the particle velocity.

Let me split those up just for organization and put the linear terms on the left and put the nonlinear terms on the right. So $F(U)$ equals some function of $G(U^2)$ on the right-hand side. At this point we can forget that we started with the Poisson equation and say, in general, this will be the form of the nonlinear equations we will have. We will keep only quadratic terms, for the most part, and we can separate into linear and quadratic.

The first cut ordinarily, when you cannot find an analytic solution, you cannot find that trick to break the problem open, is to solve the problem by a perturbation scheme and the perturbation scheme is based on the assumption that the nonlinear effects are small corrections to the linear solution and we presume we can solve the linear problem.

What you do is you introduce a perturbation expansion of this sort, where you take the wave variable (in this case, particle velocity U) and expand it in terms of an arbitrary small parameter that is a constant. We will use epsilon for this small parameter and the appropriate ordering parameter, 9 times out of 10, in nonlinear acoustics is the acoustic Mach number, say, particle velocity over sound speed, that dimensionless quantity. As I mentioned, even for the 150 dB wave, epsilon is 0.01; it is certainly small. What you want is to have the solution, each successive term be order of magnitude smaller than the previous one.

Now, if you take this expansion and stuff it into both sides of the equation at the top here, recognize that epsilon is an arbitrary parameter, you then, for a solution of this equation, must have term-by-term equivalence for every order of epsilon all the way down because of epsilon being arbitrary.

What you find, then, is U_1 . The first term is a solution of the linear problem (hopefully you know it or you would not be going this direction normally). Then you find at order of epsilon squared U_2 is a function of U_1^2 , but you knew this from the previous step. So you take this known solution, stuff it in here, so you have solved a homogeneous equation at this step, which is linear. Now you have got an inhomogeneous equation for second order, but it is still linear. The point is here, all of these are linear and you know the forcing function on the right-hand side all the way from the previous step.

Once you have set it up this way, you can imagine what is going to happen in the frequency domain. For example, suppose you have a wave at frequency omega, single-frequency source radiation. That is your linear solution. If I square sine ωT , I get a $2\omega T$ and a $0\omega T$. Let's forget about the dc for right now. That gives you a second harmonic at epsilon squared.

ϵ^3 will be the product of the U_1 - U_2 solution, so that is going to give me a 3ω . That is the first time it is going to couple back to the fundamental, so at cubic order you are going to exchange energy back to the pump wave, again. That would describe single-frequency source radiation.

What I would like to look at now is the bi-frequency case that we considered in the previous slide. In this case, I have 2 primaries, ω_a and ω_b . I square that term, $\sin \omega_a + \sin \omega_b$. I will get second harmonic of each plus the intermodulation sum and difference, so I will pop out 4 frequencies from this combination at epsilon squared. We will stop there and consider what happens in particular with the difference frequency component.

The reason is we are going to consider the case where the primary frequencies, ω and ω_b , are fairly closely spaced spectral lines. You can imagine, then, that the second harmonics in some frequencies are going to be clustered together up here and the odd man out is the difference frequency, which is going to be way down here. We want to look at the difference frequency for the moment.

[Transparency 12]

Where this gets interesting is when you have a collimated beam that is radiating this bi-frequency primary wave field. The process of generation of this difference frequency component by this mechanism, by this geometry, is called the parametric array. It was conceived by Peter Westervelt at Brown University in 1960.

Source transducer, bi-frequency waves, ω_1 , ω_2 superposed, it looks like they are beating with each other here, because they are closely spaced in the frequency domain: Well, imagine

what is going to happen physically. At any given point in that fluid the interaction of these 2 harmonics is going to be generating this difference frequency component.

Suppose it generates a difference frequency component at this point in the fluid. Think about it now in the context of, say, Green's functions. Remember, we have a linear solution of an inhomogeneous equation here. The inhomogeneous part is the forcing function from the primary waves.

The primary wave has launched a difference frequency from this point, which will now propagate as a free wave. The primary waves are still propagating, and guess what? They are propagating at the same speed and they are continuously cranking out second harmonic. The second harmonic, the difference frequency wave that is being cranked out by the primaries, is in phase with all the difference frequencies that it had pumped out before it.

Effectively what we have done here is created an end-fire array. It is a generation of difference frequency sound along the extent and axis of this beam phased at the speed of sound. The difference frequency is not generated at the transducer itself, it is generated at the volume of fluid in front of the transducer and is phased as an end-fire array at the speed of sound. From antenna theory you can do that integral and you will find that your directivity of this system -- well, we will get to that in the next slide here.

[Transparency 13]

A quick overview of what is happening here. Here is our source bi-frequency signal, 2 spectral lines. It propagates, we have a cluster here of the sum, and second harmonic signals up top, and the difference frequency way down here.

Remember absorption. As I mentioned, it is a low-pass filter. Suppose the primary waves are an order of magnitude higher in frequency than the difference, a 10:1 frequency ratio here. Viscous absorption goes normally as omega squared; therefore, the absorption coefficient is 2 orders of magnitude different between these two spectral lines. It is 100 times higher up here than here.

As this wave propagates, primary waves start tailing off due to the fraction absorption, shock formation, et cetera. In the meantime, in the nearfield here they will be generating these high-frequency components, also the difference frequency component -- less efficiently, actually. What happens because of the low-pass filtering of the fluid, these guys get snuffed out early and it is the difference frequency that is the lone survivor, far from the source.

Although it is generated relatively inefficiently -- in fact, damned poorly (it is less than a percent conversion efficiency getting it from the primaries to the difference on your best of days) -- nevertheless, it is what you are going to hear out here. It will be your higher sound pressure level.

MR. SMITH: Is that a linear scale at the RS?

DR. HAMILTON: It is Penn State's ARL graphics office from the 1970s. I have no idea. I do not know what it was intended to be. The way this is falling off, I imagine this is dB here, the way that is falling off on this side from the absorption. Otherwise, it is going to go exponential attenuation like this.

[Transparency 14]

Let's look at this thing schematically. The directivity function associated with this radiation process we call the Westervelt directivity. This is what his paper was all about,

deriving this. Here are the assumptions. The assumption is that any real beam will have a nominally collimated nearfield and then a spherically spreading farfield and then some transition region here.

To make the analysis simple, Westervelt assumed that absorption effectively terminates the nonlinear interaction region within the collimated nearfield, so you can just assume plane waves in a tube, essentially, for the purpose of doing the integral. That is what he did, primary wave ω_1 , ω_2 absorption here. Take those terms, stick them into the right-hand side of the equation and solve with, effectively, a Green's function approach; you just integrate along that line, because now you know the strength of the generation of these difference frequencies all the way along that line.

What you find, then, is a directivity function for an exponentially tapered end-fire array given by this function. What is interesting about this function is, first of all, it does not depend on source dimension. The source radius is not in here, the source frequency is not in here. That is because the source is not radiating the difference frequency.

The characteristic parameters, the so-called ka of this aperture, are the absorption coefficient and the difference frequency wave number. Why the absorption coefficient? The higher the absorption of the primary waves, the faster that primary wave field is attenuated and the shorter the array, so that is controlling the length of this radiation region. Your aperture length is 1 over your absorption coefficient, sort of.

The other interesting feature here is no side lobes in this ideal case. Why? Because in this model that aperture is going out to infinity, it is exponentially tapered out to infinity, and it is just the right way such that you are not going to get a side lobe out of it.

You can calculate a half-power angle here associated with it, and it turns out to be proportional difference frequency wavelength over this absorption length, 1 over the absorption of the primaries. The difference frequency that is being generated has an absorption coefficient, too, and when you do the Green's function integral, it is in there, it is part of the Green's function.

[Transparency 15]

Does it work? Yes. These are experiments done at Lake Travis down in Austin by Tom Muir in his doctoral work in the 1970s, primary wave in water, 418 kHz, second primary, 482 kHz, generating a difference frequency of 60-something kHz. Take a look at the primary beam directivities here, lots of side lobes on the upper primary and lower primary, Bessel function directivities nominally describe these for a side lobe down on the order of 15-20 dB, et cetera.

Here you look at the difference frequency directivity. The novel aspect is this, that the beam of the difference frequency looks very similar to the main lobe of the primary beams. This is unusual -- well, this is virtually impossible -- in linear theory, especially constrained to the limitation that your source has a given aperture radius A .

If the same source of radius A were radiating directly at the difference frequency (remember, your directivity is going as ka , that is your fundamental parameter for your beam width), this 60-kHz difference frequency is a factor of 7 or so down from these. That means your beam will be 7 times wider. You have collapsed your beam by basically an order of magnitude by generating it this way. Here is one of the benefits. Yes, you lost on the efficiency end, but you are picking up on the directivity end.

[Inaudible question]

DR. HAMILTON: It revolutionized nonlinear acoustics in the 1960s. The Navy thought it was going to be the be all and end all of sonar. That is where a lot of the initial investigations came. They liked the long-range propagation, they liked the directivity of it, they liked the side-lobe suppression of it. They hated the efficiency.

Where it is used in the Navy to a large extent now is, because of the directivity, for underwater acoustics applications and, because of the low frequency and the directivity, is for sub-bottom profiling. It is very useful there because you get large penetration depths under the sediment down to the bedrock.

Commercially, I think the biggest commercial production of this device is for fisheries, fish detection, because it covers certain frequency bands and directivities that are useful, and several German firms make parametric arrays for this purpose. That is where the money is.

What I am going to lead to here is an application that they are considering in San Diego for a home audio- speaker based on this that I want to say a few words about, where I am dubious.

This will be, I guess, a good time to take a break and I can then move into the next topic here.

DR. HAMILTON: The next extension of the parametric array is, well, what happens if, instead of a bi-frequency source, in which case you have an envelope that is beating in time with itself, you have just a pulse with a slowly varying envelope.

[Transparency 16]

A guy named Berkta at ISVR came up with a solution of this problem just a year or two after Westervelt came out with his paper; he saw the generalization to a narrow-band source, not necessarily bi-frequency. Really, the mathematics is identical, basically just transform the Westervelt directivity-type solution into the time domain.

The upshot is if the waveform at the source is some envelope E that varies slowly on a time scale compared with the carrier frequency ω_0 , it looks like these pulses down there. Then, on the axis of the beam, outside of that absorption length I told you about that effectively ends your parametric array, the pressure you will measure will be 2 time derivatives of the square of the envelope. The carrier wave disappears and you pick up 2 derivatives.

You can imagine the square of this envelope looks about like this. The first derivative will give you a positive spike here, a negative spike here, and the second derivative gives you a doublet at either end -- it is a little more smeared out when you get to this Gaussian pulse.

One of these derivatives comes from the nonlinear wave equation itself. The second derivative is linear; it comes from going from the nearfield to the farfield of a diffracting beam (that is where your JK comes in and all your linear solutions). In other words, if this were a plane wave, it would be the same solution, but 1 time derivative here, but because it is a beam, you have 2 time derivatives.

[Transparency 17 – unavailable at time of printing]

First of all, does it work? Here is an experiment we did in Austin in glycerine (this was proved in the 1960s but I will just use some data we had for another purpose). This is a several megahertz pulse in glycerine, absorption length on the order of a centimeter, 1 over your α length. This is the Berkta solution down here, essentially, but you can model the thing all the way through. Indeed, right down to the wire this is a good approximation of what is happening to this signal.

There has been, over the years, discussion about can you use this effect, if it is not difference frequencies you are excited about, could you transmit speech or other information using this parametric array in circumstances where you would like a narrow-beam transmitter or possibly you like the low-frequency aspect of it for the propagation, but mainly the narrow beam.

This was looked at in the 1980s -- in fact, it was reported first at the nonlinear acoustics symposium in 1983 in Kobe by the Japanese, by one of their large telecommunications firms -- and they experimented with what they called an audio spotlight. The idea was they made a source like about this big, generated primary frequencies in the ultrasonic range, namely, 30-40-50 kHz, and were trying to down-shift their voice signal to this demodulated band and have a propagated narrow beam.

The applications they talked about in that paper were use in tunnels for warning systems, as communication systems in tunnel systems there, and also for use in perhaps exhibit areas where you do not want cross talk between booths and you want speech to be localized.

DR. MAYNARD: What is sigma?

DR. HAMILTON: I guess that is relative to the Rayleigh distance.

They worked on that for a while and nothing particularly came of it. There were a couple of papers on it. I heard one of their demonstrations in the 1980s on it. It sounded terrible and they admitted it. They had a vocalist, I remember, and they were playing it and it was just hugely distorted. This was the big problem they were fighting, was the distortion. Also, the levels were not high enough to where they were happy.

The reason for this flip in the order of the slides is because I want to say a few things about a company that has come up with a new idea that has been shown at, at least the Penn State meeting, if not another one, for a demonstration at the Acoustical Society about a year ago, called American Technology Corporation. Woodie Norris and his collaborators have revived this idea of a parametric array in air and they are trying to package and promote this for audio applications. You see chatter all over the Internet on this right now, which can be rather entertaining reading.

[Transparency 18]

Let me describe what their device is and what some of the problems are that they may run into with it. You do not have these slides. We are giving a paper on this in Seattle, so I just pulled these out of my talk for next week. There are 3 of these and I promise, Logan, I will send these in.

Remember I told you that the parametric array directivity is determined by its length and, therefore, the absorption. For example, if you wanted to design this parametric array in air, where would you pick the frequencies? The idea is you want them ultrasonic and you want them relatively low in the ultrasonic range to keep the absorption down so your array length goes up.

I guess, Hank, you will recognize these curves. We coded in the equations from your recent paper on relaxation in air and took 1 over that and just plotted them as our absorption length for the array.

Suppose you are operating around 40 kHz, relative humidity on the order of 50% or so, you are going to have an array length on the order of several meters at the audio-frequency wavelength. Let's start from there, and that is what they have done.

[Transparency 19 - unavailable at time of printing]

Here is their audio spotlight, just schematically what they have done. The device they demonstrated has a disk, an annulus, outer radius 7 cm, inner radius 4 cm, and it has got about 60 bimorph sources around there that are resonant at around 40 kHz. They brought this device down to Texas this last semester for one of our seminars, so we walked around it with some hand-held sound-level meters. We did not get this measurement, but they claim their sound-pressure level at the carrier frequency is on the order of 140 dB and we were measuring levels on the order of 80 dB at their difference frequency.

Let's take a look at the parameters of this device based on the specs that they are running it at. Primary frequency on the order of 40 kHz -- just for discussion, let's think of it as a parametric array driving the 1-kHz audio signal. The absorption from the previous slide gives us about a dB/m at 40 kHz. That sets your array length, so this gives you several meters of distance for the aperture.

If you look at the diffraction length, the nearfield collimation distance of the beam, it is on the order of a meter, too, at the primary frequency, so these are of about the same order, so the Westervelt model should not be totally out of bounds here (at least we can use it for a first cut at the directivity).

If you calculate the half-power angle from the equation I gave you at the difference frequency, based on these parameters, 3 dB to 3 dB, you will find out that it is 17°. Well, 17, is that, good, bad, or indifferent? Compare it to what would happen if you took this source and radiated directly at that difference frequency.

Again, your ka determines your radiation pattern from this disk. Let's suppose, for convenience, it is just a circular uniform disk with no hole in it. In any case, your ka at the difference frequency for the outer radius is about 1. Well, you are not going to get even a null in that beam pattern; it is just going to come off almost like a spherical wave, so there is virtually no directivity at the difference frequency, were you to radiate directly out here, and you get a fairly "beamy" radiation at the difference frequency from this device using this principle.

And indeed, it was impressive. He got up there at Penn State in a room a little bigger than this, I guess. He would point it around the room and, sure enough, you would hear that low-frequency sound reflect off walls and he would scan the audience with it, and it is an odd sensation with someone holding a source, relatively small, to hear those low frequencies sweeping past you directionally the way it did.

We walked around the beam in Austin and you could get 50 m away walking through that beam and it was pretty focused; they were able to dump a reasonable amount of energy into that band.

The rub here is distortion. Remember, the Berklay model tells you, you are going to square that envelope and differentiate it twice. That squaring is fine if you are looking at just a single frequency, but if you are looking at different frequencies, a complex signal, you are going to introduce distortions; a highly nonlinear operation you are dealing with here. When they play speech, it sounds okay, but if you force them to put a music C.D. in there, it does not sound so good and they were reluctant to do that -- it is quite garbled.

[Transparency 20]

The purpose of the paper was to talk about that up in Seattle, what if you had a way -- and this is what they want to do -- to pre-distort. You say, okay, I know what the Berklay solution is,

I am going to take 2 derivatives of an envelope squared and that is what is going to come out due to my nonlinear acoustic process.

Suppose I try to work the problem backward and pre-distort? Let me integrate twice to get rid of the derivatives and then take the square root and now stuff the envelope in and modulate my carrier and let it go, so when it demodulates, it is back to where I started from, so that is the logic they would like to go for.

The problem here is if I have a band up here at the primary wave frequency, at the carrier, if I have an envelope and now I square this frequency or take the square root of it, it is a highly nonlinear operation, and I am going to generate harmonics all the way up to infinity, it is a huge set of nonlinear harmonics that will go along with this. The bandwidth of the transducer itself is going to limit your ability to even pre-distort the signal in the way that the parametric array wants to demodulate it on its own.

The reason this device works is it is highly resonant, and this is their big selling pitch. It is a very highly resonant, high-Q device, and you need that to dump the energy in the fluid to get those primaries up to 140-150 dB and let the fluid do the rest. If you want to get high fidelity out of this thing, you are going to have to spread that bandwidth out considerably and that is going to reduce your efficiency, and they do not have a lot to burn there, because the amplitude is not very loud, to begin with.

[Transparency 21]

If you run the numbers on that, which is what my co-authors did, just to give you an idea of what we are talking about here, if your transducer has a relative bandwidth of, say, 10% -- so we go up this line -- and we have a modulation index of about a half, and that is about where they are operating, 1 being full scale and amplitude modulation, you are looking at about 5% total harmonic distortion just because your transducer has a finite bandwidth and you are dumping all those harmonics off, even though you have attempted to mathematically undo the signal before it went in. If you have an infinite bandwidth transducer, all of these curves will go to 0 harmonic distortion in the context of that model.

Our opinion on this is it seems reasonable for speech when you do not need high amplitudes, but there is a lot of emphasis on this for -- they are trying to sell it to Sony, for example, to put it on computers and PC's. They have got Bell Labs, or Lucent, now making tests on it, and so forth. It seems right now, unless something really clever is down the pike here to come up for the signal-processing end of it, that they are up against a wall for high fidelity if they want to use it for audio applications, and that has been their Holy Grail, at least in the literature. Basically, they talk about hi-fi sound in your home, in which you can move spotlights of sound around your living room, and so forth -- you should see the P.R. in the magazines on this now (they have gotten a lot of press on it).

DR. ATCHLEY: How does the conversion efficiency depend on the difference frequency?

DR. HAMILTON: It goes as the square, so you have the omega-squared dependence.

DR. ATCHLEY: Well, no, I mean if you are trying to use it and you have a really high frequency or a low frequency.

DR. HAMILTON: Those 2 time derivatives in that Berkay solution is your omega squared at the difference frequency, so there is your filter right there. That is what we are compensating for with the 2 integrals there. That is the easy part. It is the square root that is

eating your lunch, because it is so nonlinear you cannot have a wide-enough bandwidth source to get that signal out there so that the fluid can collapse it back for you.

DR. HARGROVE: Mark, they are doing amplitude modulation of one of the primaries?

DR. HAMILTON: Yes.

DR. HARGROVE: Would they be better off doing frequency modulation?

DR. HAMILTON: Conceivably, and you could get a Berkay-type solution for that.

DR. HARGROVE: Since the conversion depends on the frequency difference, could you use that?

DR. HAMILTON: You still have the envelope squared in there, you cannot get away from that. Your derivatives would change a bit, but you still have the square root to deal with.

MR. POULSEN: This is just a concern, just a comment. The audio range will destroy your hearing? And now how do you know, if you have a dog in your house? (Laughter)

DR. HAMILTON: Yes, and now they are also interested on the effects of ultrasound on pets all of a sudden -- human beings, too -- but, yes, they are concerned about that, too. The question is, what does high-intensity sound do to your ear in the ultrasonic range? There is not a lot done with that. Certain people say, well, you cannot hear it, so how can it hurt you? (Laughter)

DR. GARRETT: The objection I have -- my picture of this is fairly simple, the length of the array is set by the attenuation length in the fluid, as you pointed out. The conversion efficiency is proportional to the number of wavelengths of the difference frequency that fit in there and the directivity.

If you are talking about expanding the audio band, you are talking about having something that is 10,000 wavelengths long, very directional and very efficient, or relatively efficient at the high frequencies, and directional, and if you are going down into the bass your array may be 1 or 2 wavelengths long, in which case you are no better off than with a direct radiator and the directionality is omni. Is that fair? Is there any way around that?

DR. HAMILTON: That is completely fair and, again, to the omega-squared dependence; that is your amplitude factor there. Your directivity is a different issue. Remember I told you one of these derivatives came from diffraction? That is on axis. If I move off axis, I lose that derivative. You are losing an omega as you are moving off axis to boot, so where are you in the beam? There is another filter, it is a spatial filter, if you are going to get picky about it.

DR. GARRETT: So it is worse.

DR. HAMILTON: It is different, at any rate.

DR. BASS: Mark, when this company came by, one of the things they left you with was the impression that their Holy Grail, their home audio thing, is probably naive in part, but it is a beautiful classic demonstration.

DR. HAMILTON: It is excellent.

DR. BASS: If we could get hold of some of these to demonstrate parametric arrays in our acoustics labs, it would certainly be nice, if there is any way we could get access.

DR. KEOLIAN: I may show that tomorrow. With all this wonderful interest, I will go and see if I can find it. (Laughter)

DR. HAMILTON: Sorry to set you up like this.

I did not mean to go here, but I thought it was worth it, since it is being bandied about a bit. Any more questions on that?

MR. WAKELAND: The thing that impressed me the most about their demonstration was when the guy was shining it over on the wall and then it sounded like it came from over there. It is probably a very small niche application, but all you wanted it to do was provide for you that kind of directivity through reflection. Does that take care of some of your distortion problems if you use only high frequencies and the low-frequency part came from some --

DR. HAMILTON: The narrower the band you go with this in terms of the frequency, you get rid of the problems that Steve talked about and the diffraction, so all of these start collapsing as you go down to one spectral line. I suppose you can boost a note or two with it, but with the kinds of standards that they put on audio systems, I just cannot imagine, if they are doing harmonic distortion tests on any of this -- I am not familiar with how they calibrate these THX's and all this kind of stuff or what they do for just marketing -- you are right, I think it is niche applications and I just have a hard time imagining yet that music is going to play a part in that for a while.

MR. HARDIMAN: Is there a requirement to phase steer at all?

DR. HAMILTON: Only in the sense that in a linear array the primary beam itself could be phase-steered. To the extent you can steer the primary beams you could steer the difference frequency.

[Transparency 22 -- unavailable at time of printing]

I will wind up the parametric array by saying you can turn this thing around backward and you can use it as a receiver, just to show you it can be done. Here the applications are quite remote, but you can take one high-frequency beam and pass it through a low-frequency beam, so the high frequency is now your pump and you are trying to pick up a low frequency at some signal frequency ω_s , which is way down here on the axis.

As soon as you interact this spectral line at ω_0 with your signal at ω_s , you will generate these sum-and-difference frequency side bands. What you are doing is that side band is generated by the propagation of this signal from the source to the hydrophone here and effectively, again, you have created an end-fire array of length L , but you have done it with just 2 elements. The continuous elements of your end-fire array are now the fluid elements themselves in-between.

DR. SABATIER: Have you ever made a pulsed echo system?

DR. HAMILTON: Yes. Again, this comes back to doing things like measuring B/A. You can run pulses through each other and look at the phase shift of one by another and try to back out of that -- nonlinear properties.

DR. SABATIER: There is a transmitter and a receiver?

DR. HAMILTON: Yes, is the short answer.

[Transparency 23 -- unavailable at time of printing]

This was Jim Truchard's doctoral thesis, "Nonlinear Acoustics: Parametric Arrays." Does anybody use Natural Instruments Labview? He is the founder of the company and the president. I hope he remembers his roots. These are some beam patterns from his dissertation. These do have side lobes, because it is a finite array of length L , so you do not have that exponential taper any more. These were, again, done at ARL, Texas, up at Lake Travis.

One picture I have seen of an application or a deployment, or at least a test run, of this, was along a submarine. They used the span of the submarine for the length of the array, put the pump on one end, the source on one end, the receiver at the other, and shot the beam down the axis of the sub, so the whole length of the sub was the aperture length. Of course, you have to steer the sub to sweep through.

On the other hand, you have a very directional -- you get a narrow-band receiver there in space for that low frequency, whereas the cross section of the sub, for example, is much smaller. They have run tests like that, but to what extent it has really been deployed, I do not know.

[Transparency 24]

We are back to the main notes now, Section 2, on what happens when you get a shock. Your world does not fall apart completely. There are actually some simple patches you can use here to put the pieces back together, again, literally.

As this wave propagates, eventually a peak will want to overtake a trough. Unlike water waves at the beach, sound waves cannot break and they start catching up. The gradient and the fluid increases, the shears increase, the heat conduction increases, the losses increase, and the upshot is it is filtered back. The wave tries to steepen forward, viscosity and heat conduction try to keep it from doing that, and you can hit a balance.

If the wave is very, very strong, very high amplitude, compared to these viscous effects, then you can maintain on the scale of the other time scales in the waveform a very thin shock that, if you do not care for the moment about the structure of that shock, like what is its width, you can get some very simple analytic expressions to describe the propagation of that shock in the wave. This is called weak shock theory.

The shock wave relations in their exact form, the Rankine-Hugoniot relations, were derived at the end of the 19th century, but for use in acoustics it was not until the 1940s and 1950s that weak shock theory was really developed by the likes of Hans Betha at Los Alamos and Landau and Witham, and names like that, who were the ones who really put the pieces together so that you could keep things together.

If you go back to the exact shock relations of Rankine and Hugoniot, basically they said, okay, I know all of my fluid and energy variables on one side of the shock, I know it on the other, I know I have to conserve mass, I have to conserve momentum, I have to conserve energy -- energy not just in the sense of the acoustical but internal due to the shock passing over that fluid -- and if you insist that those qualities must balance across the shock, you get what are called jump conditions relating the variables across that shock.

They are very nonlinear equations and there are very few analytic solutions of this, but if you now look at the acoustic limit, where the jump quantities are small, then you find that the velocity of the shock has a very simple analytic expression. It is in reference to the sound speed C_0 of a linear sound wave, $\beta/2$ times the average of the particle velocities just ahead of the shock and just behind.

Notice that if U_a equals U_b , in other words, that shock jump goes to 0, we have twice the particle velocity divided by 2, we have C_0 plus βU , which is our solution that we had previously for the continuous segments of the wave.

The point is, this is consistent at that level of approximation with the earlier equations we had. What do you do? You use the Poisson solution to describe all continuous parts of the

wave. You get up to this point, you now know the amplitude of the wave just ahead of and just behind the shock; therefore, you know how fast that shock is going and you can now describe the propagation of that shock within the waveform itself and you have been able to combine the continuous segments of the wave using this coupling relation.

That is what weak shock theory is all about and it was put on a firm footing especially by Blackstock's papers in the early 1960s, showing how powerful this method can be and really formalizing it into a kind of algorithm approach.

Landau came up with a very clever alternative view of this weak shock relation. It turns out, if you take that Poisson solution I mentioned before and let the wave steepen over on itself, let that waveform go multivalued, then the place you put the shock is at that point where the line would cut off equal areas of that solution that you had previously allowed, and now erroneous, to go multivalued. This is a direct consequence of this relation, so this gives you a geometric way of sticking the shock in your waveform and, for certain problems, can be a convenient way of doing things.

One subtle point here, subtle for us these days, because we ignore it, but it drove the 19th century physicists up the wall and basically stopped them from solving the problem -- I kind of slid this right by you, it is easy to forget about -- is look what we are doing here. What I am telling you to do now is go back and use that lossless gas solution for the continuous segments of the wave and stick a shock in-between to couple them together.

What does a shock do? You are dumping energy into that shock, you are losing energy, it is irreversible, from the wave. What is going on here? A couple of things can happen, and this is what bothered them. One of the things that bothered them was, well, there is a discontinuity of the wave. The continuous segments of the wave from behind the shock catch up from the rear and those from ahead fall back into it.

From the wave point of view, this thing has hit an impedance mismatch. It will reflect. These wavelets will actually reflect off this shock and you no longer have a progressive wave from the instant in time that shock is formed in the fluid, so that wave is rolling along, rolling along, and all of a sudden the shock forms. It is going to start sending signals back.

These are internal reflections from shocks and then your whole progressive wave equations fall apart in an exact sense. They did not like that. The other thing they did not like was energy. They said how can you use lossless theory, because you have losses going on here and you have got an energy drop going across that shock, so how can I be taping them together with theory that has no losses?

The answer to both of these questions was very subtle and was not worked out until the 1940s and the 1950s and that is why it took until then for this weak shock theory to come about. What you find is that the jump in the energy, the jump in the entropy, and the jump in the Raman [phonetic] invariance, which tells you whether the waves are going forward or backward, are all third order in acoustical variables, and our theory is stopping at second order. If those variables kicked in at second order, we would be hosed and you could not use this.

It is a very subtle point, but crucial as the underpinnings for this entire theory. You do not see viscosity in any of these equations, yet you are properly accounting for the energy loss in this wave. It is kind of incredible.

MR. WAKELAND: A couple of questions. You said the shock forms and all of a sudden you have an impedance mismatch and you start getting reflections. Why all of a sudden? I mean, don't you start getting some as it forms?

DR. HAMILTON: No, the equations of motion emit a perfect isolated progressive wave going from left to right, provided it is continuous. The equations of motion emit that exactly.

MR. WAKELAND: Can you go through what the symbols in the $V_s = C_0$ equation there mean?

DR. HAMILTON: C_0 is the propagation speed -- this is a good time to advertise. There is a great book that just came out on nonlinear acoustics here. I highly recommend it. It goes through a lot of this stuff in detail. A lot of the figures I am just ripping out of here. So buy many. (Laughter)

This is the position of the shock in your XT plane. This is the velocity of the shock as a function of time. C_0 is your small signal sound speed, it is a linear one. γ_0 over ρ_0 for a perfect gas, square root. That is pretty much all of them. U_b is the solution of the Poisson equation right at that point, one point before it fell into the shock. U_a is the solution in the limit as we are sneaking up to the shock from the other side, and those are the points that we use to evaluate that quantity, which tells us how fast the shock goes.

MR. WAKELAND: The particle speed, the pressure?

DR. HAMILTON: Particle velocity, fluid particle velocity, $P/\rho C$, roughly.

MR. WAKELAND: And beta is -- ?

DR. HAMILTON: The coefficient of nonlinearity.

MR. WAKELAND: Right, okay, so that is all positive?

DR. HAMILTON: Yes.

MR. WAKELAND: U_b is bigger than U_a , so the whole thing --

DR. HAMILTON: Not necessarily. All of these were positive, but this axis could have gone through here.

MR. WAKELAND: I guess what I am asking is, is V_s bigger than C_0 ?

DR. HAMILTON: Yes, that has to be, that is right.

Beta is always positive, so the wave can steepen only this way and shock.

DR. HARGROVE: Beta is not always positive. You just said so earlier.

DR. HAMILTON: The wave always steepens forward.

DR. HARGROVE: No, it does not always steepen forward.

DR. HAMILTON: For beta positive. For all the fluids I have ever worked with, beta has steepened forward. (Laughter)

DR. HARGROVE: The fact that quartz has an anomalous sine in various thicknesses of plates in shock tubes to tune the shock thickness, because you can increase rather than decrease the shock thickness at the end of the shock tube by putting it through something that "undistorts" the wave and thickens the shock thickness.

DR. HAMILTON: So they undo it down the line.

MR. WAKELAND: When you have the airplane flying over and you have the cone coming off the nose and you calculate the angle, you have got to use a speed that is higher than the ordinary speed of sound in the fluid, is that true? Or you have to use something else?

DR. HAMILTON: This projectile is going supersonically. The wave is going to come off of that Mach angle, but it is going to propagate nominally at C_0 , and it will be an N wave in space, and the head shock will look like that and the tail shock will be negative, because it is a negative wave. It will look like the letter N.

The back side of a shock is always higher in amplitude than the bottom of the shock and that is what keeps the shock maintained. You have to have an overpressure driving it. Does that help?

MR. WAKELAND: I still just do not know VS is. I thought it was the speed at which the shock was moving along.

DR. HAMILTON: It is. It always -- I take it back, it will not always. I wish I had a pen here.

DR. HARGROVE: I think having gotten into it in this much detail, it is worth pointing out that the shock is forming in the vicinity of where there was the zero crossing and the sine wave, which we usually think of as being at regular intervals, but you move the zero crossing, so to speak, now so that the positive part of your wave and the negative part of your wave will no longer be of the same time duration.

DR. HAMILTON: This answers your question exactly. Here is the N wave, here is curve 1, curve 2. It starts out like this. What is going to happen? This pressure is higher than this pressure; therefore, the overpressure will drive the shock forward. It will move forward, but because these amplitudes from behind are catching up with it, the amplitude will go down.

Relative to this zero crossing here, the shock will go ahead and be decayed. Second, the reverse happens on this side. The shock will be maintained because there is still an overpressure here. The overpressure is required to maintain the shock; however, the mean velocity here is less than equilibrium and, therefore, it will slide back.

Here the shock goes forward. Here the shock goes back, but it will have its top cut off the same way. What will happen is the period will increase as square root of distance and the amplitude will go down as square root of distance.

MR. WAKELAND: The center of the N is moving at C_0 at a slightly higher speed because it is spreading out.

DR. HAMILTON: That is right. Thanks, that helps.

DR. GARRETT: I have a question. Hank, didn't the Army at one time use the width of that wave to try to do ranging of artillery?

DR. BASS: The length of the pulse?

DR. GARRETT: Yes.

DR. BASS: I do not think they have any artillery systems that are based on that.

DR. HAMILTON: We can get an exact analytic solution for an N wave. In fact, the slide before was kind of cluttered. It was from Rudenko's book. He shows you how to do the equal area rule to get at that and you can turn this example into an N wave, if you want. The point is, however you do it, it is simple.

I did not write the full expression, I regret now, but the exact solution for the pressure is instead of X it is $1+BX$, let's call it, and the period goes as square root $1+BX$. It is the same factor. The point about B is it has amplitude in it, source amplitude, not current amplitude, and it has beta in it. The point is, all the others are material constants.

You actually have an absolute measure of what your source amplitude is if you have a perfect N wave propagating, or a perfect spherical wave (you can turn this into spherical wave equations). By curve fitting you can figure exactly what the amplitude of your source had to be by working backward, by watching the rate of decay of this.

The point, again, here is that it is not square root X, it is square root of -- I took the high distance limit, thinking of this asymptotically -- but B depends on your source amplitude, P_0 , let's call it, at the source. Plotting this function carefully you can do that, and Blackstock does that to calibrate his high-frequency condenser microphones, because he has no references to use. He actually uses these N-wave equations, and there is one of his N waves from a spark source out of his lab. He propagates these things, curve fits them backward, knows all the other parameters from just material properties, and figures out what the amplitude is and the frequency response. Very clever. It is absolute. One measurement, bang, you have it, you have got the whole spectrum.

MR. WAKELAND: So he knows how far away?

DR. HAMILTON: Yes, he takes several measurements and curve fits that curve, that is right.

MR. WAKELAND: So if you know how loud the thing was each time, if it were always the same amplitude, you could figure out the --

DR. HAMILTON: Yes, right, that is right.

DR. HOFER: Mark, that is free propagation? Not in a tube or anything?

DR. HAMILTON: He does it with spherical waves. All this plane-wave stuff I am giving you, even with the weak- shock theory, there is a simple transformation to get you to cylindrical and spherical waves. It is no extra math -- one-stop shopping.

[Transparency 25]

I mentioned renewed interest in N waves. Sonic booms, because of high-speed aircraft, supersonic aircraft, that NASA was looking at. Lithotripsy is another biggie -- I will just throw it up here in case you have not seen what it is. It is used to focus shock waves to disintegrate kidney stones in your body and they are pulverized to grain size where you can pass them more pleasantly -- I am told. (Laughter)

I actually saw one of the trucks in Austin. It is just like a big moving van, it just comes right up to your house and they just do it in your driveway. (Laughter)

They just put you in, drop you in the tub. Here the person is immersed in it, but I like this one here. Here it is a dry application where they are coupled with a bladder here. You put a spark plug in water at one focus of an ellipsoid and blow the sucker off about 1000 times and you try not to sink it with heart rates and things like that. The shocks come in and you hope the stone stays where you think it is -- a good application of phase conjugation here.

How come the water bag on the guy's belly? Here is an acoustics 101 question: How come the water bag on the belly?

[Transparency 26]

Let me show you what the waveform you are hitting him with is. You are hitting him with this shock, about 100 MPa at the focus. Why the water bag? You are hitting, from the point of view of the sound wave, you are propagating through a liquid, the body, and you are seeing air, you are looking at a pressure-release surface.

That thing is going to turn around as a rarefaction wave and you could get -- well, if it were a solid, you would get spallation. What was happening, they found, were big bruises on people's backs or on the opposite side they were coming in from, from rupturing of blood vessels, so they put a layer of water so the reflection does not occur at your interface, it occurs at the bag's interface.

DR. SABATIER: But that is only a factor of 2 and what happens to all the rest of the waves there and on the way back?

DR. LAUTERBORN: It cavitates them.

DR. HAMILTON: At the interface?

DR. LAUTERBORN: Yes.

DR. SABATIER: Does it not hemorrhage? You are suggesting it hemorrhages at the surface, but somehow it does not hemorrhage all the way along?

DR. HAMILTON: Oh, you do, you cavitate in there. In fact, they think the cavitation at the stone is the principal cause mechanism for the stone breakup. There was an argument for years about is it cavitation or is it spallation -- is it cavitation, I guess you would say.

DR. LAUTERBORN: At the moment they think it is cavitation, because the breaking of the stone depends on the frequency that you use. If you wait too long between 2 pulses, nothing happens, and it is spallation that should be independent, so people think it is bubbles left and these bubbles are compressed, so you make continuously cavitation and that destroys the stone.

DR. ATCHLEY: There is tissue damage around it, but not as much as a knife slicing through your body.

DR. LAUTERBORN: The tissue is more flexible than the stone and, therefore, it --

DR. HAMILTON: Try not to get stones. (Laughter)

MR. POESE: Why is it just kidney stones? Why not gallstones?

DR. HAMILTON: They thought it would work for gallstones, too, but I guess the material properties of the stones, for whatever reason, the brittleness of them, are not such that this lends itself to --

MR. LAUTERBORN: Kidney stones are more brittle. They are like what you find in water, very tough, and so they would not break up so easily.

DR. HAMILTON: My guess is there are more gallstones than kidney stones, so that was a bigger market they were hoping for.

There are three main manufacturers of lithotripters, Seimens, Dornier, and there is a third. Dornier is, I guess, the widest used. It is the spark plug in an ellipsoid. One is piezo. The deal with that is it has a much tighter focus because you can control the beam much more. What I have heard, when Larry Crum talks about it, is that, although in principle that seems nicer, you can focus the energy more, that presumes you hit the target, and I guess they would rather have a weaker beam, broader area, and not be boring holes through the subject in a shotgun approach.

I guess the third one is an electric hydraulic coil.

[Transparency 27]

The beauty of the weak-shock theory is you can now stick this weak-shock result into your analytic solutions that you have derived from these Poisson solutions, and the classic solution here -- I am just going to show it to you, you can write for a plane wave, and the solution is used

very frequently, so I want to give it to you -- you can write the propagation of a plane wave starting at a single frequency in this form.

Here σ is a dimensionless distance relative to a shock distance; in other words, $\sigma=1$; it means that the sinusoid at that frequency has just formed a vertical tangent. In the pre-shock region you have this Bessel function solution derived by Fubini as a graduate student in 1935, and here you have the sawtooth solution. For σ greater than 3, you again have an approximate analytic solution. The fact that you are going off $1/N$ in frequency means you have a sawtooth wave looking like this. In the middle you are in kind of a no-man's land here from 1 to 3, but Blackstone gives you an integral you can do to get yourself from here to here.

The point is you have the whole darned thing. This is possibly the most popular or widely used analytic explicit solution in nonlinear acoustics, both the Fubini and the sawtooth.

Let me pause here for a second and note that with this X bar that I am normalizing by is your shock-formation distance, and let me just say a word about that. Look what it depends on. This is for a sinusoid, source pressure P_0 , frequency ω , coefficient of nonlinearity β .

If I double the source pressure, I have the shock-formation distance; that is fairly obvious, at least that it would go in that direction. How come if I double the frequency I also have the shock-formation distance? Why should it depend on frequency? Peaks and troughs are half as close when everything else has been maintained the same, so you have to go only half as far for these points to catch up to each other in a relative sense as they propagate.

[Transparency 28]

I will show you an experiment here, which is interesting. Let's look at the sawtooth solution for a moment. Something rather fascinating occurs. Let me write it here. This is the solution: Particle velocity U launched by a source, amplitude U_0 at frequency ω . Here it is. Now envision this experiment.

σ is the only dimensionless parameter that varies here, barring the time, for any given spectral component. Suppose my receiver is at a fixed location X . I fix my source frequency and β and C_0 are material parameters. The only thing I am going to vary is source amplitude U_0 . I am just going to crank it up and watch what happens to my level over here.

Let's first look at what happens to the solution. This means that I am letting σ get large eventually, large compared to 1. I can throw away the unity and I now write the full expression down here. Look what happens. Source amplitude U_0 , appearing here, cancels with that here, and I have an expression for the sawtooth wave that is independent of the amplitude of the source. U_0 no longer appears in the problem.

Here is the experiment graphically. I have a source at a fixed location, receiver at a fixed location. My x axis is amplitude on this end. My y axis is amplitude, let's say, at the fundamental (I am just going to look at the source frequency level, you can look at any of them, but let's just look at that).

For linear theory, 60 BN, 60 B _____. For nonlinear I start coming around this curve; the σ gets large compared to 1 and eventually flattens out and I know exactly where. I cannot get the level any higher at that point in space. How can this be? Where is it going?

It is coming out of the source. It is a perfect source, it is not an electronic problem, nothing like that.

MR. POESE: You are heating up the fluid?

DR. HAMILTON: I am heating up the fluid, exactly. You are dumping the energy at the shocks. It just so happens that for a sawtooth wave of that form, doubling the amplitude halves the shock-formation distance and the ratios work out just right, such that it is a breakeven at the point out here in the fluid.

You do not get that for the N wave, you never saturate. Instead of going up linearly, you end up going up square root; that is the equivalent of saturation, but you never flatten out with an N wave. It has to be a sawtooth.

MR. POESE: You do not see this in a standing wave, either, apparently, in a resonator?

DR. HAMILTON: You do, sort of, and I will say a little bit about that.

MR. POESE: Because these macrosonics guys seem to be able to beat this, they do not have this saturation problem.

DR. HAMILTON: In the next half I will show you.

[Transparency 29]

Let me show you an experiment and then we will go out for a break here. These are the types of experiments that one typically does, plane-wave tube, driver at one end, anechoic, termination at the other, so we to some extent assure some type of progressive wave motion here. This is an experiment done by Blackstock in the 1970s, just verifying that, indeed, these saturation curves do happen on length scales of the order of 15-20 m in air at audio frequencies from half to, say, 3 kHz, so it is not hard to get there.

Certainly even in underwater acoustics and ultrasonics at megahertz frequencies it is very easy to hit these saturation levels, and it happens to all of your harmonics. Indeed, this is a problem. Oftentimes you would like to offset this saturation suppression, you want to get that energy farther out there. In particular, there is a standing wave application for an acoustic pump that I will talk about later, what the inventors did to beat the physics here.

Let's go on a break now and I will talk about that when we come back.

DR. HAMILTON: Once you get to either complicated waveforms, for example, random waveforms, noise, or you start running into odd dissipation -- or any dissipation, for that matter, but particularly nonclassical ones, that is, nonthermal viscous, such as relaxation or tube-wall attenuations and waveguides and such and a little bit of dispersion kicks in -- you do not have a hope of getting an analytic solution of the problem.

[Transparency 30]

You go to numerical methods and there are several very simple codes you can get a lot of mileage out of and can pretty much run on personal computers, desktop computers.

This was one of the early ones that was particularly successful and is still used today for some of the sonic boom work. It was developed by Mike Pestorius and David Blackstock as Mike's doctoral dissertation. In fact, it is referred to in the literature as the Pestorius algorithm.

It is very simple and it cuts to the core of the fundamental principle behind most of these numerical methods, and that is this: If you are in a coordinate system that is traveling with the wave, the things that are happening to you as you are riding along with it are happening very slowly, over many wavelengths.

That permits you to decouple these effects. Numerical methods have various terms like "split step algorithms," and so forth. The idea is this. You propagate over a step and you take into account one effect. You bring your solution back, pick up the next, bring your solution

back, pick up the next, and now go and do it again for the next step. You can do this even in nonlinear problems. The fact that your operator is nonlinear does not really have much to do with that, provided you have a small enough step size; you can guarantee some kind of convergence.

One of the first that properly accounted for shocks and arbitrary absorption and dispersion was this algorithm. The way it was done was simply using the equations I just gave you, really no more than you have seen on the viewgraphs, to put this whole code together.

Take an arbitrary waveform, come in here, and propagate according to the weak-shock routine. That is the formalism I just gave you. Either we are going to go over step $\bullet x$, take the continuous parts of the wave and steepen them, if you have a shock, recognize it and move it forward according to that weak shock relation. Now we are at $\bullet x$.

Let's now move back, because we may not have taken -- well, we have not taken into account the effect of absorption and dispersion on those continuous parts of the wave, because that weak shock theory was lossless everywhere but at the shock. How do we get that in there?

The way you do that is you now come back, transform into the frequency domain and pick up your imaginary and real parts of your attenuation coefficients to give you that absorption and dispersion to bring you up to here, which is done at this step here. Now bring yourself back into the time domain and around and around you go, you propagate your way down.

You can do all of this in the time domain, provided you have operators to give you this absorption dispersion, and we have codes that do that, too. These codes run very fast, especially in 1-dimensional cases, literally seconds. They just fly.

Again, time domain steps, Poisson solution for continuous segments of the wave, weak shock theory for the discontinuities in the wave, flip into the frequency domain and throw in a complex attenuation term over the step $\bullet x$. That can be an arbitrary law for the absorption and dispersion properties of the fluid.

[Transparency 31 – unavailable at time of printing]

This works quite well. I will jump all the way ahead to this noise slide -- there are a couple of examples there and cartoons of how shocks eat up shocks, but you can see it in this experiment. This is an experiment that Pestorius did. This is the reason he wrote the code, is to work this problem of the propagation of finite amplitude noise.

The experimental waveform here was used simultaneously as the input to the code here. There were ports drilled into that waveguide I showed you at various distances out to about 85 feet, then the measured waveforms were compared to what the propagation routine predicted. It is rather incredible to have noise, nonlinear, with all of these effects modeled so well.

This simple piecewise building block approach to dissecting the evolution of a wave indeed works over sufficiently small length scales. These step sizes, he has gone out to 85 feet, were not particularly small; as I remember, probably half-foot increments or so. I mean, you are hopping your way out there and you can still account for these effects.

Notice what happens here in terms of these shock speeds.

[Transparency 32]

I will go back to the earlier slide here, the stick figure. Take, for example, this cartoon. This shock has a mean velocity that is higher than this shock. This has an overpressure relative to this, this shock will catch up with this one. As this shock catches up, eventually it will

overtake and merge, and this is irreversible. Once a shock catches a shock, the 2 shocks become 1 forever after.

An interesting process takes place here. As the wave starts out, it steepens up. What happens when a wave steepens? The gradients go up, the spectral contents go out, and the frequency spectrum is broadened. You are pumping energy up to the high-frequency end.

Now, if you have a random waveform with all kinds of different shocks in it of different sizes, eventually shocks will start merging together, because they are propagating at different speeds. As the shocks merge, the spatial structure of this waveform now decreases, so the Fourier content slides back down. Initially you get an up shift of energy and then you get a down shift of energy for the propagation of noise. This is one mechanism that contributes to that.

[Transparency 33]

In this slide you can actually follow one of these guides in-between. Here is a little bitty shock. Here it is, here it is, and it got eaten up here and it is gone. Exactly the same thing is modeled over here on the right-hand side for the computations.

[Transparency 34]

Here is the frequency spectrum corresponding to that experiment, input spectrum, say, at a foot, centered around 1500 Hz. As it propagates out, 2 things happen. You are generating high frequencies -- these are the sum frequencies and second harmonics, if you think of these as just a bunch of spectral lines interacting with each other. The spectral lines will also generate difference frequencies, and that is what is pushing up your spectrum down here.

As the wave propagates further, absorption kicks in and starts peeling off the higher end, so from 49 feet to 73 feet you have dropped 510 dB because of the viscosity, the energy loss effects at the high-frequency end.

David told me that he was going to ask graphics to take off that little spike there. He thought, oh, that is silly, let's get rid of it, and they did not, they left it in there. Do you know what it is, you waveguide people? It is the first non-planar mode cutting on.

[Transparency 35]

You can do interesting things, also. Instead of the propagation of pure noise, you can look at the interaction of a tone with noise. This was of interest for aircraft noise, strong tonal components coming out of an engine interacting with a broad background. This is numerical simulation.

Suppose you have this high-amplitude tone, 150 dB, interacting with a spectrum of noise at 110 dB (again, these are SPLs in air, read 20 μ Pa). If that high-amplitude tone were not there, this spectrum at all distances would be the same, it would not matter, absorption aside, it would just propagate.

[Transparency 36]

Because this tone is here, it changes the wave nonlinearly. This spectral line will interact with all the other spectral lines contained in this band and generate effectively some frequencies. If you think about the perturbation scheme I told you about, of course, you get first order, you generate second harmonics, then third harmonics, fundamental the second generates the third, and so forth. It staircases its way out here just like this.

As it propagates, eventually it just flushes this noise all the way up into the high-frequency spectrum, so it takes that low-frequency band of noise and runs it all over the frequency domain

along with its own harmonics. This was in exact correspondence with the experiment that they performed, Webster and Blackstock, 1978.

[return to Transparency 35]

Here is the tone propagating in the absence of noise, 150 dB, I guess, these harmonics -- it is right at the source, so primarily coming out of the source itself. We burn out a lot of drivers down in Austin (JBL loves us). As this propagates, it is generating its set of harmonics according to the Fubini solution I showed you earlier.

Here is the band of noise without the tone. As it propagates, nothing really happens. Now put the 2 in together and you can see what happens. That tone, just as we saw in the numerical model, runs all the noise up to the higher frequencies, so you have a complete spectral redistribution of this low-amplitude noise band taking place simply by the presence of a finite amplitude single tonal component at the source.

You can turn this problem around a bit, and I think Bob is going to do an experiment on this, and Anthony, I know, has published a paper on this, you can use noise to -- this is, in a sense, absorption -- absorb a tone, so you can work it the other way. You can use noise of finite amplitude to suppress a tone that is embedded in it. That is another kind of dual problem of this. We will not go into that here.

Let's change gears for a second. I need to show you the full gore for a few slides here, just what the full equations look like, so you can appreciate what we have thrown away.

MR. WAKELAND: I got the impression when I was reading these beforehand that -- when you say finite amplitude, you mean not infinitesimal?

DR. HAMILTON: Yes. I am so used to it now that I forgot. I never liked it either when I first heard it. Finite, right, as opposed to 0, not finite in comparison to infinity. Infinitesimal, yes, thank you, good point.

[Transparency 37]

These are the full equations of motion for fluid that you really need to really do the problems fully. They are 3-dimensional. They take into account shear viscosity, bulk viscosity, and heat conduction. These are just copied right out of Landau and Lifschitz, who have an excellent discussion of their derivation.

Just a few words about these to show you some of the terms we are going to throw away right off the bat. The idea is eventually we are going to keep only quadratic terms, so right off the bat this whole third term here and the energy tensor leaves, because it is quadratic in the field variable, plus we pick up a small quantity in the shear viscosity here. That is a very small term. This is small for the same reason.

This term often goes out -- notice it is linear, actually it is a linear loss term, linear and it has got a loss coefficient here, shear and bulk viscosity -- this is your vorticity term. This is very important in acoustical streaming problems.

In potential flow problems, which is the bulk of what we do for free space propagation, it is not important. It is not important mainly when you are away from boundaries, on the order of a wavelength or so away, because once your flow starts out your rotational in potential, it will remain potential and this term will not kick in; it is separate and apart from everything else that is going on in the fluid.

If you strike a surface, all bets are off. You may need to account for this vorticity term there.

MR. POESE: The fluid has momentum, right, and that takes energy, and that is how it is a loss?

DR. HAMILTON: Right, yes. There is momentum loss to the fluid and that is what acoustical streaming, which is D.C. flow, is, precisely that. The wave has been absorbed and therefore it lost momentum. The fluid needs to carry that momentum somehow and the way it chooses to do it is it flows. That is acoustical streaming. If you shut off the absorption, you shut off the streaming.

There is radiation pressure without absorption. Lighthill has a paper saying there is streaming without absorption. I do not want to go there, it is very subtle. I can show you where to look for that if you care.

MR. WAKELAND: So the streaming is not really a nonlinear phenomenon?

DR. HAMILTON: Streaming is a nonlinear phenomenon. It varies quadratically. This is one of your linear terms in the streaming equation, but there are others there, too. I have to time average these, so I do have quadratic terms. This is a linear term. I start a new rotational flow, it will remain so.

DR. GARRETT: How does Michael get around circulation theorem? If you get streaming without viscosity, then you change the curls in the liquid and for that there has to be viscosity.

DR. HAMILTON: I repeat --

DR. GARRETT: Okay, okay. (Laughter)

DR. HOFER: Maybe this is not quite such a subtle question, but do you have streaming in what is called quartz wind [phonetic] where you have an ultrasonic beam?

DR. HAMILTON: Yes.

DR. HOFER: Is that because you are losing energy out of the beam that is not necessarily dissipative or --

DR. HAMILTON: It is dissipative. As soon as α goes to 0, it shuts off -- except in Michael's paper. (Laughter)

It is a very subtle point having to do with edges. It is a non-1D flow and he gets into some funny things there that I never really understood or cared to.

MR. HARRIS: If you have a randomly heterogeneous media, can you go that linear term?

DR. HAMILTON: Yes, thank you, this is for homogeneous fluids. Everything is different for homogeneity.

MR. HARRIS: The little turbulent areas of particles and things like that, can you still ignore that linear term?

DR. HAMILTON: Oh, no, because then you would have interaction with the sphericals, yes, and that is, in fact, a loss mechanism. In this acoustic resonator we will try to get to here, they are driving the amplitudes in the standing wave so hard that they are making the flow turbulent, so the acoustical energy as a loss is going into turbulence, which is not immediately into viscous. It will get there, but there is another stage in-between. That is a good point.

Another assumption here: All of these loss coefficients -- this is your thermal heat conduction, shear and bulk viscosity -- are assumed to be independent of not only position but

amplitude of the wave. That is not exactly true, either, but that is pretty good to take at face value.

[Transparency 38]

Let me quickly show you the model equations that kick in here. If I throw out all of the losses, I can get a wave equation in terms of one acoustical variable. Let me show it to you, because it is classical. If you choose the velocity potential, the scalar potential, we have thrown out the rotational part, this is the full deal.

Your linear operator is on the left and you have these 4 terms on the right: 2 are quadratic, 2 are cubic in the velocity potential. That is exact for a perfect gas, no losses, except it is rarely used in acoustics. You cannot solve it except by perturbation and so forth and, in any case, it is valid only for a perfect gas, so if you want to do liquids, you are in trouble right from the get go and cubic terms are of no use to you and what about viscosity?

But purists like talking about this, so just to show you, there it is.

[Transparency 39]

The practical way to go here is to say really what terms are important and how can I consider a broad class of fluids. Here we get some guidance from the way Lighthill looked at the problem -- he did some of the pioneering work in the 1950s on derivation of the Burger's equation, which I will show you.

The idea is this. You recognize there are 2 small parameters that characterize the smallness of the terms in your equations. One is your acoustic Mach number, which we talked about, and the other is these loss coefficients. So what you say is I will retain only those terms that go up to second order in terms of these combinations, so my order epsilon terms are my lossless linear terms -- that is my lossless linear wave equation that you can make careers on.

At second order I have lossless nonlinear effects of quadratic order. Those are the ones I have been talking to you about up to this point. I also have linear terms with loss coefficients in front of them. Those are my viscosity and heat-conduction terms, but I will classify these as also second order.

I am considering both of these small corrections to the linear solution. I will not go to third order here. That is the rank ordering. Throw out these and keep these and these. In these terms an additional corollary is I can use an order epsilon relation and substitute it into either of these, because the error in doing that is third order, so you can do a lot of simplification really fast. Those are the ground rules; just go with them and you can run amok and drive all equations based on this guidance.

[Transparency 40]

When you do that, this is what you get. Continuity, momentum, and energy equations look like this. I have written them in a special way here, where I have identified these terms in red as the Lagrangian density, which comes up again and again in nonlinear acoustics equations.

The only point I want to mention here is this -- let's take a look at it -- it is the difference in the kinetic energy density and the potential energy density. These are the same for a progressive plane wave. L is 0 for a progressive plane wave. Those terms disappear.

Look what happens. Your momentum equation becomes linear; that is what I was mentioning before. Nonlinearity comes primarily in fluids from your continuity from your

continuity equation and state equation for progressive plane waves. You can get a long way by just forgetting about the nonlinearity and the momentum equations.

Your energy and state equations are combined into one here and written here.

[Transparency 41]

I just want to touch on these. If you take those equations and collapse them into one, you get this equation, first presented in this form by a Norwegian nonlinear acoustics group in the early 1980s, Aononsen, et al. Your lossless linear wave equation -- here. All your viscosity terms are lumped into one linear term called delta. Delta is referred to by Lighthill as the sound diffusivity. It has your bulk and shear viscosity and your thermal conductivity, so it all comes into one parameter here; you cannot separate them at this order of approximation, you get them all together.

You have 2 nonlinear terms. One is kind of like a monopole, they talk about source term, it is a p^2 . The other is distributed in space and it involves the Lagrangian density. As I mentioned for plane waves, you can throw that Lagrangian term out and you need only the first 3.

This equation, that top equation, is through second order in the sense that I have defined it, complete. It is not restricted at all to progressive wave, spherical wave, anything. It can be compound waves. They need not be 1-dimensional, they can be multi-dimensional. It takes into account shear/bulk viscosity, and heat conduction. It is rare, even if you go to higher order than this, that you can identify those effects and measure them. That is about as far as you would want to go for a general wave equation. Everything from there distills from this equation from my perspective.

The Westervelt equation is simply gotten by throwing out the last term. These are the equations you are going to see all over the literature. That is why I want to present these, as just kind of a menu. The so-called Westervelt equation is this equation with the L term thrown out, so it works very well for progressive waves. It does not work for standing waves, because what is the condition of throwing L out? or waves that are sufficiently directional that they look like plane waves? A directional sound beam fits that very well.

[Transparency 42]

The final step I want to take is to derive the classical model equation of nonlinear acoustics and that is the Burger's equation. Here you use this scaling procedure, again in a more subtle way, but in a very powerful way and in a way that has broken open methodologies for how to look at nonlinear wave evolution processes. This was pioneered to a large extent by Khokhlov and his colleagues at Moscow University in the 1960s.

Let's start with the Westervelt equation and say I care only about plane waves going that way, but I want to keep account of the nonlinearity, I want to keep account of the dissipation, also. Now, in your minds let's rationalize their approach.

If I throw out nonlinearity, $\beta=0$, I know the simple plane wave solution to this problem. You open up Kinsler and Frye, pull this out, there is your thermal viscous attenuation coefficient. Let's look at the other limit of this equation. Let's throw out the viscosity term and keep the nonlinear term. We know from our Poisson solution that we looked at earlier that we darned well better recover it, and we know what that solution is for a sine wave. It looks like this.

Look at the arguments of these functions. In both of them I have lumped the time together in a retarded time frame τ , $T-X/C_0$. I am moving with the wave, so I have put myself in a

moving coordinate system. The only other functional dependence in this problem is through space.

Remember our ordering scheme. Both of these depend on X as that small parameter times X . In this case it is pressure times distance. In this case it is loss terms times distance. But functionally speaking, it is the same concept according to that ordering procedure, small parameter times length scale. This is the scale at which these effects change in the wave, because it is a very small parameter times your coordinate. It is called the slow scale.

You recognize explicitly now this slow scale with epsilon and say my solutions and equations must have this functional form. Do the coordinate transformation, throw out all cubic order terms, just do that religiously, and out pops the Burger's equation. You will find that these are the only 3 terms that matter.

First of all, when you expand your X derivative you will not that the last term is epsilon squared, but remember you are multiplying times the pressure. That puts it third order. Out it goes. You have reduced your operator to first order in propagation direction. That tells you *a priori* it is going to be for only progressive plane waves.

Progressive plane waves. Here is your viscosity term. Here is your nonlinearity term. This is the simplest model equation that completely describes the combined effects of dissipation and nonlinearity on progressive waves.

MR. POESE: Say again how you know it is only for progressive plane waves.

DR. HAMILTON: You need a second-order operator to have compound waveforms and you have only first order in your space derivative.

What is remarkable about this equation is that it has an exact solution. These guys, Hopf and Cole, came up with it. You can, through a nonlinear transformation, turn this into the diffusion equation, and we know what that solution is, we have an analytic integral for that, many analytic closed form solutions for it, so you can get that solution. The problem, of course, is bringing it back, but you can do that analytically in a number of cases. It is rarely done, but it is there, so it is intriguing from a mathematical point of view.

MR. WAKELAND: That very last factor on the page, could you put parentheses in there somewhere? Is that P^2 and then you take the first derivative of it?

DR. HAMILTON: P^2 , then take the derivative.

[Transparency 43]

This equation is really easy to solve by computer. We give this as a homework problem, frankly. Now you are off to the races, because you can do all kinds of things very simply. Throw the equation in dimensionless form, space derivative, 2 time derivatives and your viscosity -- here is the Goldberg number that Steve mentioned -- and here is your nonlinear term, $pdp/d\tau$, which is the way I rewrote that.

This is in a dimensionless form that is very convenient, because sigma is scaled to your shock formation distance, so everything here is order 1, all your variables.

Assume you can expand this pressure as a Fourier series. Take this, stuff it in here, and you will find that you now have coupled equations for the spectral amplitudes. You can blaze through these equations on a simple Runge-Kutta routine. What is great about this is now you are sitting in the frequency domain. You can now forget that you started with the Burger's equation, forget that it was thermoviscous viscosity with quadratic frequency dependence and

make that coefficient complex and throw in whatever absorption, whatever dispersion you want, because you are solving this with complex numbers anyway. That is no problem.

It is a very, very efficient code for most problems, even for well-defined shocks. If it is single-frequency radiation, you need only a couple hundred harmonics. Often you get away with a couple dozen.

DR. MAYNARD: Where do you define X bar?

DR. HAMILTON: X bar is your plane wave shock --

DR. MAYNARD: But where is it defined in our books?

DR. HAMILTON: On the page with the Fubini solution.

This shows you, also, there is one parameter telling you the characteristics of these solutions. There is only one free parameter in this equation: We call it gamma. Blackstock labeled this the Goldberg number in an early 1960s paper. He originally called it G . It is the ratio of nonlinear effects to absorptive effects through here.

The reviewer of the paper balked and said that gives too much credit to this guy Goldberg in this paper in the 1950s, he never understood the implications of this number to the extent you are giving him credit, do not call it G . So Blackstock renamed it gamma. (Laughter) The reviewer said, "No problem." (Laughter)

In his paper it is gamma everywhere in the text. In the figures it is crossed out, it is G 's, because he could not change the figures back then. Back then they were all hand-drawn and he did not want to go there. At any rate, it is a little bit of history.

[Transparency 44]

I did not plan on covering this, but because we are doing model equations, one other approach to solving nonlinear wave problems is because of Jay Maynard's talk the other night on Hamiltonian formalism, there is another whole approach you can take here and I just want to give you a snapshot of that, because we have used this with tremendous success for surface waves.

I want to show you just very briefly how it can be done using Hamiltonian mechanics. The problem I want to show you is surface waves, a surface wave in a solid being a Rayleigh wave, Scholte, or Stoneley wave. These are waves that decay exponentially with depth in the interior of the material.

If the solid is anisotropic, is a crystal, instead of simple exponential decay into the solid, you get oscillatory decay into the solid. The hardest part of this problem often, in solving the linear problem, is getting these eigenfunctions for these surface waves in these crystals.

[Transparency 45]

The question is how would you solve this system for the nonlinear case? I am going to give you the overview for this, but this would work for any acoustical system, plane waves, what-have-you. The point is, you will see how general it is.

You expand your wave function, your particle velocity here, in terms of eigenfunction ψ . These are solutions of your linear problem. What we are going to do now is assume that your nonlinear solution is a slow modulation of this linear solution and solve for these coefficients a . Watch how this is done.

[Transparency 46]

You construct your Hamiltonian, calculate your kinetic energy, and you will find it is proportional to these amplitude coefficients a . You now look at these amplitudes and say I am going to declare those to be my generalized coordinates. Once I have done that, by definition my generalized momenta become the derivatives of kinetic energy with respect to these velocities.

Then you have Hamiltonian canonical equations at this stage and your system is closed. In principle, you can solve the problem from here by eliminating variables once you have an expression for the potential energy V . That is exactly what you do, express the potential energy V and, therefore, the Hamiltonian in terms of these amplitudes.

It is really algebra after that and you get coupled spectral equations just as I showed you for the Burger's equation.

[Transparency 47]

This is what your potential energy looks like in general for an arbitrary medium. It is expanded in different terms, because it depends on the constitutive nature of the problem. The V_2 's give you the linear theory. They are quadratic in order, they are quadratic in the strains. Here is your elastic strain tensor that we saw yesterday with your coefficient c that Al was measuring.

If it is piezoelectric, here is your electric field contribution and here is the coupling between electric and your strain. There is your piezoelectric. All of these are linear effects. They are quadratic in the wave variables for the energy, they are linear in terms of the dynamical equations.

Cubic terms: This is where your nonlinearity comes in. You get cubic in the strain, cubic in the electric, and then you get cross-terms between the electric and the strain. The point is, this is completely general for an arbitrary medium with arbitrary piezoelectric effects involved in it.

For the surface wave and for most plane waves you can take all of these derivatives. Your plane waves, if you are differentiating in this direction, are just bringing down factors of ik in your propagation number. For your surface waves, they have wave numbers going in these directions, but they may be complex. Nevertheless, you are still differentiating exponentials. This whole thing can be done analytically and the integral can be done analytically, you get an analytic expression for the whole shebang.

[Transparency 48]

For the surface wave you have a remarkably elegant and simple solution of this problem looking just like that Burger's equation I showed you in this frequency domain. Principally, you can take this and stick it on your same computer and solve it with your Runge-Kutta method.

The only difference here is you have a nonlinearity matrix, not a constant coefficient sitting here, and that you calculate. You know that in terms of your material constants and eigenvalues and you are off to the races and you can run it. We have gotten tremendous mileage from using this equation. This was derived by Zobolovskaya, who was one of Khokholov's students who worked on the Burger's equation originally.

We have applied this to all shapes and manners of surface waves and I will show you just one of these to show you what the propagation looks like. In fact, I will go straight to the experiment.

[Transparency 49]

This is an experiment that was done in Germany by Lomonosov and Hess generating nonlinear Rayleigh waves in fused quartz. A laser beam incident on the sample -- this is for isotropic solids (I will not go into the more complicated cases) -- launches a surface wave of high amplitude, and 2 photodetectors. Two reflected beams are used to probe the signal, so it is non-contact source generation, non-contact measurement. The trick here was being able to get photodetectors of high enough bandwidth to capture the shocks in these waves. That is the experiment.

Probe is blasted here, sample is blasted here, first probe beam at 2 mm, second probe beam at 18 mm. The name of the game is if you give us the waveform at station 1, can we match the waveform at station 2, just like the noise experiment I showed you earlier.

The shock-formation distance for these waves is on the order of a millimeter and the frequencies are on the order of 50 MHz, very fast shock formation here, very high amplitudes. It is a highly nonlinear system. The theory works like a charm.

[Transparency 50]

Here is the waveform that is measured at station 1 at 2.3 mm, vertical component of the Rayleigh wave, horizontal component of the Rayleigh wave, frequency spectrum. We took this waveform, wrote it as a Fourier series, stuck it in the code, propagated it with Runge-Kutta, pulled it out the other end, and stuck it on their measurements at 18 mm: bang-on agreement, especially these spikes at the shocks. These are unusual attributes of surface waves that are characteristic of surface waves.

When a surface wave steepens, it does not form a sawtooth; it generates spikes around the shocks that can be very, very big -- we have not shown that case right here, but this is one of the distinctive features of nonlinear Rayleigh waves that is predicted by this theory, and we have measured these aspects.

DR. GARRETT: Do you have any way of explaining why these things cut up that way?

DR. HAMILTON: Rayleigh waves are special in that the nonlinearity is non-local. In that sound wave I showed you, if I know the pressure at that point in the wave, I know how fast that part of the wave is distorting. If I know everything I need to know about one point on the Rayleigh wave, I do not know what the nonlinearity is. I have to know everything about the wave over the entire time or space axis.

How come? How can this possibly be? Imagine a pulse of a Rayleigh wave on the surface, a square pulse limited in time. Remember what a Rayleigh wave does. Its depth penetration is proportional to wavelength, so as I go beneath the surface, that wave is actually spread out in Fourier space, if I do a spatial Fourier decomposition of that, because I am stripping off those high frequencies.

Waves that do not look coupled at the surface are coupled in the volume of the material. This gives rise to a nonlinear operation that gives you the cut. That is as close as I can get to it. It is a very complicated process, but they are very special in that way.

MS. ZOU: How do you get your vertical component and the horizontal component?

DR. HAMILTON: From beam deflection, laser beam coming in and the angle tilts coming off. It turns out that that tilt in angle is proportional to the vertical component of the surface deflection.

DR. QUATE: Does this argument hold for plate modes?

DR. HAMILTON: Insofar as surface modes are part of that, yes. But if I took a waveguide and got rid of the elastic coupling, no, it would not.

[Transparency 51]

I want to wind up by jumping ahead to dispersion. I want to show you something special about that. I am going to jump all the way to this slide on waveguides as just a kind of warm-up for this. What sets nonlinear acoustics apart from nonlinear optics is one word: dispersion. Optics has lots of it in media that are required for the nonlinearity that make it dispersive. Acoustics has very little of it.

What does that mean? It means in acoustics -- remember the parametric array? -- everything is going at the same speed, I am traveling along with some wave pumping energy into another, they are phase-locked, and that is why I shock. In optics those waves are going off at different speeds. They do not phase lock unless you do special things to make them do so and then you are coupling only certain spectral lines. Dispersion is knocking the phase speeds out of whack for all the other spectral components you might want to interact with.

Where do you get big dispersion in fluids? The 2 examples that we worked through in the book -- I will go through only one here -- are geometrical dispersion due to waveguides, because you get tremendous dispersion in waveguide modes, and bubbly liquids; huge dispersion and huge nonlinearity due to bubbles. We will not have time to go into that.

Let me talk about the waveguides, because it pertains to a resonator problem I want to get into. Just to show you what happens with dispersion, we did this experiment at U.T. Suppose I generate two waves in this mode, a plane wave mode to 001 frequency and a higher wave mode in the 10. These 2 are going at different speeds down the waveguide, ω_1 and ω_2 ; again, I will get sum-and-difference frequency generation.

The primary waves are moving at different speeds and so will the difference frequency and sum it is trying to generate, so everything is going at different speeds here. If you think about it, an analog in mechanical systems is driving a mass-spring oscillator off resonance. ω_0 K/M is its natural frequency.

I am now going to start driving that mass at ω different from ω_0 . The particular solution will beat with the homogeneous solution and this thing will oscillate in time, it will beat before the transients die out. The exact analog here is happening in space. Now we will beat in space as the waves are in phase, out of phase, in phase, out of phase, but it is a nonlinear beating, it is not a linear effect.

[Transparency 52 -- unavailable at time of printing]

We measured these oscillations in this waveguide and it is very simple. Just look at the wave-number mismatching and you will know exactly what that oscillation period is. These are the same types of things you would find in nonlinear optics.

The point here that I want to get to is the stronger the dispersion -- in fact, for example, the dispersion for the difference frequency is higher than that for the sum frequency. The coupling is worse, the periods are faster, and if the dispersion length (we call it) is shorter, the amplitudes get less and less, because it can get to only such a high amplitude before the phase knocks it out again. In other words, at very high dispersion this thing will just be bouncing along down here in the noise level.

As the dispersion goes up, the oscillations come in, the amplitude goes down, and what are you doing? You are blocking the amount of energy that is getting into that spectral line is the other way of looking at it. You are preventing energy from getting to that frequency component.

You can use dispersion to keep energy out of frequencies you might want. Why would you want to do that? Suppose I had a fundamental wave -- remember, as it propagates, it is going to generate a second harmonic line, third, fourth, fifth, and so forth. What if I did not want that wave to shock?

If I put this in a dispersive medium that was so dispersive that I cannot get the second harmonic going, that energy will never get out of there. You will stay in that spectral line, you will prevent shock, or offset or postpone, at least, shock formation to the extent you can suppress reactively the amount of energy that is coming into those components.

[Transparency 53]

This applies also to standing waves. I want to wind up by talking about a hot topic in Virginia, in Richmond, a company started up by Timothy Lucas, called MacroSonix, which has gotten extraordinary press, beginning at the San Diego meeting, where they unveiled this to the public.

I do not have the numbers on this (I put this together last night, so I am kind of remembering some of the details of this), but I will give you the big picture, explain what they are doing, and why it works, because you will be hearing about this.

The idea is to make a pump, a fluid pump, by exploiting the resonance of sound waves in a chamber. The original motivation was to avoid the use of CFCs in mechanical piston-generated compression cycles. The idea is this: You put the fluid in the chamber and you drive it with a shaker and drive this cavity in its fundamental mode.

Now what you do, if you look at the end here and you look at the pressure oscillation going negative/positive, negative/positive, you can now tap these valves and phase them in such a way that you are drawing fluid through there. You are actually able to do a pumping cycle here by timing these valves. This is what his idea was.

There are few moving parts. You are driving this thing at several hundred hertz, I believe. The resonators are about this big. It is far closer to solid state, certainly, than the pump that you have got strapped onto the back of your refrigerators, and because it does not use seals, and so forth, you do not have to worry about the alternatives to CFCs that are now being proposed.

The problem was lack of dispersion, initially. If you just put this thing in a Coke can, in a cylinder, and start shaking this thing, the fundamental, which fits into this mode, the second harmonic will fit exactly at half-wavelength into a mode and the third harmonic are all singing, they are all in resonance here.

Thinking about it as a plane wave, if I am a sound wave propagating in this direction, I am just going to keep bouncing back and forth. Well, all my harmonics are in phase with me, because I am bouncing off rigid end caps here. Each time I come back and forth, back and forth, I am resonantly pumping these harmonics and I can shock, and you will shock. You do not get sawtooth waves, you get squirrely kinds of waves, but they have got shocks all over the place.

They had unbelievable amplitudes in here. They were getting up to Mach numbers of 1. I mean, it ceases to be acoustics; you are taking a slug of gas and throwing it against that wall and then throwing it back this way, again, so it is hardly the acoustic approximation any more.

The problem is they were saturating. Even in a standing wave you saturate and you could not get to the amplitudes they needed because of that nonlinear saturation I talked about to get the problem done. How do you get rid of that? Dispersion.

[Transparency 54]

What Tim realized was if I flare my cavity, I get dispersion there and I knock that phase matching out of existence. Just the linear solution here. If you look at the Webster horn equation, this linear theory here, you can solve this for an exponential horn, a flare rate that is exponential, put in a propagation wave number that is complex, solve for it. It is a very simple basic acoustics problem. You get an incredible dispersion relation here for the phase feed.

It goes infinite near cutoff, approaches free space wave speed at high frequencies, but the point is, it is hugely dispersive. By taking this resonator and flaring it in various ways, you have thrown off the phase synchronization between all of these various harmonics and that is precisely what they have done.

Instead of those shocks, you can get other types of waveforms. They do not look sinusoidal any more, still crazy things are going on in there, but you are not saturating, where you were before, and you can get the sound amplitudes out much, much higher than you could previously.

The cute part about it is they show that by shaping the resonators in different ways you can change the shape of this waveform to your advantage for tapping the valves. For example, you want a long negative part compared to a short high part for mass transfer, because your gas is much more expanded in the negative phase and you need more time to get the fluid in, as opposed to spending a lot of time in the positive end and a little time in the negative end. You can actually play with the wave shape on top of suppressing the shock formation and this is where they are headed with that.

MR. POESE: The wave is shaped by tuning which harmonics are excited and which are not?

DR. HAMILTON: Exactly, effectively, to first order.
This is how it works. What they call resonant macrosonic synthesis, RMS, is the slogan that goes out there. This is what they are talking about.

CHAOS AND NONLINEAR BUBBLE DYNAMICS

Werner Lauterborn
Drittes Physikalisches Institut
Universität Göttingen

[Transparency 1]

DR. LAUTERBORN: I will speak on the topic of chaos and I hope I can bring that topic a little bit nearer to those who do not know it already. What we did with chaos was mainly in connection with nonlinear bubble dynamics, and I will speak on that topic.

[Transparency 2]

This is the content of what I am going to speak about. I will give a philosophical introduction, so we will gently slide into the topic, and add some historical notes, a little step further. Then we will go to the basic notions, starting with state space and these sorts of things, and go on with what maybe is at the moment at the heart of chaos, that is, nonlinear time series analysis. That is something different from Fourier analysis and correlation and all that, really something different, and I will speak about that.

Then I will come to the topic of nonlinear bubble dynamics, first, theory, how bubbles behave according to the equations we have, and then experiments. You know from the talk of Anthony Atchley that there is quite a good fit with these single bubbles, how they behave. They collapse and have all the rebounds that fit quite well.

[Transparency 3]

Because chaos is a very special subject, from time to time people like that topic. For instance, when they discovered kinetic gas theory, the very word "gas" comes from chaos. It was first called chaos, but after some time it was a little bit difficult to speak of chaos and they simply called it gas. Now we have it free again, so again we can speak of chaos. From time to time we reinvent chaos.

[Transparency 4]

If you ask a philosopher what he can say about our world, for instance, Wittgenstein came up with "The world is all that is the case." He had that in German, of course, and that is, "*Die Welt ist alles, was der Fall ist.*" I thought about that, whether that is really all we can say.

[Transparency 5]

If you ask a physicist, he has more specific words for that and from time to time he adds a little bit. A physicist, for instance, around 1900, saw the theory of physics as not really a very old subject. Around 1900 he would come up with "the world is statistical," everything was statistics.

After some time they learned that is not all, not everything is statistics; there is something more. Einstein found that out and he said everything is relative. Whatever you do, you have to add relativity to everything. That, of course, was not the end of the story.

If you go ahead, you find out that quantum mechanics appeared and people doing quantum mechanics. They say everything has to be described by quantum mechanics, also everything. You have to add that. You know the problem with relativity and quantum mechanics, that is not settled.

Now is that the end of the story or have we something else? Now you may discover why I did that. That is the next step. The next step is, we would say, the world is chaotic, and to almost everything we can add that. Whatever you think about, somehow it is nonlinear -- that is the basis of that -- and somehow everything is chaotic.

I have no idea how the story will go on, but that is the momentary status. I would say if you want to work in an area that is really challenging, then you will work in that area. No wonder, when Glynn Holt came to us and we told him that is the area and showed him that, he was converted. Indeed, he now has very nice papers on chaos.

[Transparency 6]

What may be the basis of chaos, what can we learn from it? I tried to find that out and to formulate it and I came up with what may be the main topic, the main outcome, in that area, and that is determinism does not imply predictability, at least not long term predictability.

If we know a law of nature, that means we have almost nothing won, because what we want, when we learn some law or when we find some law in physics, is that we can predict things, how they behave, predict for the future, and now we learn, with chaos theory, we cannot do that in the long term.

All we do in chaos physics is try to find out how far we can come with determinism and how far we can come with prediction, and so on. Chaos theory has very much to do with prediction, but we cannot predict, even if we have a simple deterministic law, really, what is going on.

I have an example, and that example runs as follows. We have a population of mosquitoes and I give you the law, it is a deterministic law -- it is going like that. They reproduce, of course, and we are getting more and more of them, and I just take a factor of, say, 10 per year. Of course, if you take a factor of 10 per year and you do that for a million years, all the earth is crowded with just mosquitoes. That cannot be, so there are some, say, birds or even human beings with some chemicals, that wipe that out. I say that is kept at 1 billion mosquitoes, that means 1 mos, and it is always kept below, because above that we do something about it. That is a very simple law.

If you start with that, you do not want to predict how the population goes on, and I have here a whole list. We start with a 0.526 mos and we have a factor of 10. Then we have 5.26, but then we do something about that, so we have, on the left here, only 0.26, and then we take a factor of 10, and you see how that operates. This is what, mathematically, is called a Bernoulli shift; you just shift the digits. There are many laws of nature that just shift the digits somehow, all to the front.

Then what happens is, whatever you do, whatever you measure, 6 digits, 7 digits, 8, you have no chance. The physical law shifts the digits somehow (it may be very complicated). After some time, you know nothing, so whatever you do, but you have a deterministic law. If it is of that kind, that you shift, then you cannot predict in the long run.

That is, I think, the main thing we have learned so far from chaos physics and that is the philosophical part.

[Transparency 7]

How was everything started? These are the historical notes. In any case, somehow it starts, somewhere are the first ideas, so history somehow starts and chaos comes into being.

[Transparency 8]

I tried to find out for the acoustics part how chaos came into being. As we now know, nonlinearity is essential for chaos. We look for nonlinearity. We have heard this morning something on nonlinear acoustics, so that was really the start. Mark has told us that if you have two frequencies and you have some nonlinearity, you can get the difference frequency and get what is called the combination tones -- you see that there -- that was very early.

Music: Music is a nonlinear science, because whatever you do in music, you generate sound and sound generation is always nonlinear, because if you have a linear equation, you can only do something proportional. You need the sound, an amplifier, that would be a linear amplifier, but that is not a musical instrument. Whatever you do in music, it is nonlinear.

That was first noted, that if you have two tones, you can observe the difference tone, and there are two or three persons who did that first. Tartini had the violin and Sorge had an organ. If they played two of these tones, they observed the difference tone. Later, but around the same time, they found out that there were even higher nonlinearities, but at that time they had no idea, of course. They found, for instance, what we now call the cubic difference tone, and you really can hear that.

If something is played loudly, then you hear all these tones, and that can be quite annoying. At that time there was much discussion of whether these tones are produced in the ear or whether they are produced in the medium. Today we know you can produce them in the medium and you can also produce it in the ear. It is just that you need the nonlinearity at the right place. You can have both.

[Transparency 9]

That was, I would say, almost the very beginning, but I am sure the old Greeks, maybe, knew that also, but it was not reported.

The next nonlinearity that was observed and reported was the Faraday experiment. Faraday was experimenting with his figures, where you pour sand on a plate and let it vibrate and then the sand is springing and hopping around and forms the nodes and so. He decided he would pour water on that.

He built a little instrument like this one up here and shook it up and down. He wanted to see the lines of vibration. He noted a very peculiar thing. He was vibrating that little sheet of water at a frequency f , but it was vibrating at half that frequency, the water surface, and he was really astonished.

Rayleigh repeated that experiment and also was astonished and made a theory -- Rayleigh almost immediately made a theory for some experiment. Even today, and that was in 1831, this is an experiment that you can do and that is not yet totally and fully understood.

I have a paper from 1996 in *Physical Review Letters*. If you shake a liquid up and down, you can publish that in *Physical Review Letters*. (Laughter)

It is a very, very difficult subject to measure how the surface oscillates in 2D, or maybe even 3D. That is a very, very difficult problem. It was solved in that paper that you can make this nice little view of the surface.

This is such a simple experiment. Of course, we have also done it and photographed it, and what you can observe --

[Transparency 10]

-- I brought it with me, (a Ph.D. diploma student, Christian Merkwirth, did it). If you shake this little liquid, it forms structures, for instance, like that. It depends on frequency, on viscosity, on amplitude. You can have different types of structures depending on these parameters.

In this case, a very regular structure appears by just shaking it. You do nothing else, you just shake it. You can have spirals,

[Transparency 11]

lines -- it is very peculiar -- and you can have chaos.

[Transparency 12]

If you shake it a little bit too hard, then this regular structure breaks up and you can follow the breakup of this structure and you can write down equations. It is awfully complicated to write down equations of how that could happen and how you solve them. That is what you can do by just vibrating a liquid.

There are parts in the structure that are vibrating at a never-repeating frequency, going up and down, going up and down, down and up, not at the frequency you put in, at a totally different frequency. That is really peculiar, because it is just a little bit of water that is going up and down.

[Transparency 13]

That was the experiment we did, acoustic cavitation. You put a sound field into a liquid. This is the apparatus we had. We have a piezoelectric cylinder, all piezoelectric. Dump that into the liquid. We had a hydrophone and listened to the output. You put in a single frequency. Later on in the experiment, I will tell you, we observed very peculiar things.

You put in just a simple single frequency. Then something very peculiar happens. If you photograph the cylinder (you see that below) and look along the axis, a lot of bubbles appear that form such dendritic branches. Still today it is not understood how that comes about.

Now you see that somehow bubbles are involved and, therefore, this bubble dynamics stuff came up. Nonlinear oscillators are involved. Before we go into analyzing the phenomenon, I will give a little introduction to the basic notions--

[Transparency 14]

-- that have been developed over the course of time in nonlinear dynamics to cope with these nonlinear systems. It is not just, as we heard this morning, that some wave steepens, but totally different methods.

I call them basic notions although some are such elevated notions in mathematical terms that it is really hard to grasp.

[Transparency 15]

I start with state space. It is well known as phase space from thermodynamics, but in thermodynamics, where you have these particles, your state space always is 2-, 4-, 6-, and so on, dimensional; that means it has an even dimension. State space need not be even dimensional. You can have 3 dimensions, 5, and so on, because what you are doing is, if you have a system, then you look at what describes your system and that you take as coordinates.

You take coordinates X_1, X_2, X_3 up to X_m and that may be pressure, that may be voltage, anything that you can measure. You put one variable on one axis and you take as many as you need to describe your system, and if you need 5, then you take 5 that describe your system and then you have a 5-dimensional state space.

I put that down here. What happens then is you can describe your system by the values of your, say, 5 variables, 1 V here, and that pressure there, temperature, this one, along the axis, and that makes up a point in your state space: That is this point, X_0 . It has m entries and is a vector.

Your system is not static, it evolves in time, so after some time it occupies some other point in state space. If you have a continuous dynamical system, you have an evolution equation and you can write it down like that, as I put down here [Transparency 15].

You have such an evolution equation, whether written down or not so that from an initial point, X_0 , you, during time, come to the next point in time. If you have, for instance, 5 functions, it is a whole set of 5 differential equations of first order.

μ , that is very important, you always have some parameters and these parameters also span a space, called parameter space. The parameters can alter and may occupy a d dimensional space, for instance. That is the notion of state space in nonlinear dynamics. You always look into the state space and how your system behaves in state space.

It behaves, for instance, like this red curve here, and that is called a trajectory. Your system, as it evolves in time, goes along a trajectory in your state space.

[Transparency 16]

The mathematicians have done a lot on these things. For instance, they tried to classify, if you have a system, what it is doing or what it can do if you have a set of equations. They came up with not really much can happen and invented the notion of attractor. There are only a few things that can happen and that is really astonishing.

For instance, if you take different initial conditions, your system can go down just to a simple point in state space; they call that a fixed point (for physicists it is just an equilibrium position, so every trajectory goes to that point). For instance, a pendulum that is damped goes down and then is at a fixed point, no evolution any more.

The next thing that can happen, mathematicians found out, is: your system can go round and round on a trajectory and they call that a limit cycle. It may happen that from wherever you started as an initial condition you go to this limit cycle. An example is the van der Pol oscillator or, say, any oscillator you have. If you buy a generator and you want to have a sine wave, then it is done in that way. You generate a limit cycle. Radio and T.V., all rely on this limit cycle. They produce it and then they do something with that.

The next, more complicated, thing that can happen is a little bit involved. It is called a torus. You may start somewhere in state space, then your trajectory goes to the surface of the torus and goes round and round on it. If that happens, you have 2 incommensurable frequencies in your system, that means 2 frequencies that are not related by rational numbers. So there is really something involved, because to have 2 frequencies that are not related with rational numbers, that is really something involved already. These motions (or a torus) are called quasi-periodic, because they are almost repeating, but not quite, and that may happen, that may really happen also in physical systems. The irrational numbers are really occurring in physics before chaos. That is a topic of chaotic systems and that came up in only the last, say, 20 years, because it is very difficult to calculate and to plot that by hand, so for sure you need a computer to plot that. Only with the advent of the computer could all this scale stuff could be done. Without computers, we would not have chaos, or chaos theory -- even that would not be. You need a P.C., at least, for chaos.

That [Transparency 16d] is called a strange attractor, or also a chaotic attractor. This is a very simple example from a Rössler attractor. The trajectories are going around and, if they are inside, they are going outside, then they spiral inside, then they are going outside, without ever repeating. The equation for that is just a simple set of three equations -- I have not put them down.

That is essentially what can happen; all other things are mere additions. For instance, you can have 2 attractors. You can have, in one part, everything going to a fixed point and, in another part, everything going here (to a torus), but you have mainly only these 4 possibilities, so that is very fortunate.

Now, how to plot these peculiar attractors? This [Transparency 16d] really is a projection, because for continuous dynamical systems mathematicians can show that you cannot have a strange attractor in 2 dimensions; you need at least 3 dimensions for a continuous dynamical system to have a chaotic attractor. In a plane you cannot plot it.

This is a 3-dimensional object. It starts with 3 dimensions. Of course, chaotic attractors go up to 4, 5, 6, and so on. That is the real trouble, because in 4, 5, or 6 dimensions we lose intuition, so we need very clever methods that the computer can do -- it can do what we cannot do -- to extract information from these (higher dimensional attractors).

[Transparency 17]

What you can do -- Poincaré, one of the fathers of nonlinear dynamics, has invented it -- is the following. If you have a 3-dimensional object, and these here are the trajectories, you put a plane into the 3-dimensional space and observed trajectories only when they go through that plane.

He took one point going around, got another point, going around, got another point, etc. The mathematicians have a theory for that. They call the operation a map. This point Q_1 goes under the Poincaré map to Q_2 . For instance, this point Q -- if you have a limit cycle, a periodic orbit (trajectory) -- goes to $P(Q)$, so that is the same point. It is a fixed point of the Poincaré map.

You see what happens. A limit cycle under this Poincaré map, if you make a Poincaré section, goes to a fixed point, so you go one dimension below by that, and that is really true. If you cut a torus and a point hops around, you come to what looks like a limit cycle, but the point hops around, it is not really a limit cycle.

But you have gained one dimension and you can do that with a chaotic attractor, too, and you have gained one dimension and you have the real chaotic attractor, you can see it, and you can do that to really see the structure of a chaotic attractor.

[Transparency 18]

This is a numerical calculation of this simple equation. That is what is called the double-well Duffing equation, a Duffing oscillator. The potential for that is something like that. If you have U here, say, of x , and you differentiate this curve, you end up with this one (that is the force), and if you drive that system, and you can imagine somewhere some ball going around under this forcing, if you do that with these parameters, you take as damping 0.2, as driving 0.3, as frequency 1.24, you calculate that. You can do that on a P.C., easily.

If you do the following: you take these ωt at zero, 2π , and so on as t goes, and you cut it - you always plot a point when ωt is 0, 2π , 4π , and so on, so you strobe the system. Then you

look, then you do not look all the way around, then you look again, and then you end up with exactly this plot. That is the numerical calculation.

MR. WAKELAND: So if you have like a ball in a bowl that looks like that or on a wire, a bead on a wire, and it is going back and forth and you are driving it, I do not understand why that is 3-dimensional state space.

DR. LAUTERBORN: This driving you can write as $d(\omega t)dt$. You can add an extra variable, ωt you call ϕ , for instance, $d\phi/dt$ is ω -- you can add that to the list to make it a 3-dimensional system. That is a trick, to plot time in that way together with the other variables. The state space in this case is $R_2 \times S_1$, S_1 being a circle, and that is a trick.

I will show a movie later when we come to the Lyapunov exponents where you can see what happens with this S_1 . You can go with $\omega t=0$, $\omega t=0.1$. You can cut it at different places of ωt and then you see that the state space is $R_2 \times S_1$, a very complicated thing, not an R_3 .

Yes, I forgot to plot this state space, really, but later you will see a little bit more.

This was just theory, a numerical calculation. You can also measure that. That is really astonishing and we did the experiment with a pendulum.

The pendulum obeys this equation, but these are measurements, a real pendulum that is driven, so I have no numbers for that. At the same phase ωt of the driving, a measurement was made of the angle and of the velocity, the angular velocity. That is plotted here -- maybe that is angular velocity, I forgot -- and if you do that for certain frequencies and certain amplitudes of the driving, certain parameters of your pendulum, you measure that form, 50,000 points. That goes on and on and you get that structure.

You can really measure that in experiments, and this was a very, very simple experiment. You can imagine if you do it in some not-so-simple experiment, as with bubbles, for instance, you also have that, and I will speak about that later. A chaotic attractor is a real thing. You can measure it.

That also was a diploma physicist's work: to get the experiment up, drive the pendulum, measure its motion, measure period doubling, and more.

[Transparency 19]

Now these strange attractors are real things. We have learned that. They do not come from nothing, because normally a pendulum, you know, behaves quite periodically, going up and down, back and forth. But sometimes they are like this: chaotic, and how does that come about?

You must alter a parameter, of course. Then there is a certain parameter where this change appears. First everything is regular, then a change appears and everything gets chaotic. How can that happen?

I show you again what happens. The mathematicians have found out, and that is just from the theory of differential equations. Mathematicians have found out that there are only four local bifurcations, as they call it, local because it appears at one point of a parameter. You come along, everything is regular, something happens, and that is called a bifurcation.

That is a qualitative change, not just to go with a number, 5, 5.1, or something, there is no change, but there is really a qualitative change of an attractor when a parameter is altered. Now you know there are only four different attractors, so there is only a certain number of points that you can have a change between, from a torus to a limit cycle, from a limit cycle to a torus, from a fixed point to a strange attractor, and a few more.

[Transparencies 20-21]

Mathematicians have found out there are four of these local bifurcations. One is called Hopf bifurcation. You change a parameter in your equation and then what happens is a fixed point blows up and you get a limit cycle; something qualitatively different appears.

[Transparency 20, upper part]

An example is a van der Pol oscillator. When you change a parameter, c , in these equations, you come through 0 and then suddenly oscillations start. That is exactly what you need if you want to have an oscillation. You need somewhere in your system a Hopf bifurcation. Then you get an oscillation.

Musicians know that. They have a string and they start, alter the parameters, they press along the string and then oscillations appear, so that is how it is done.

[Transparency 20, lower part]

The second bifurcation that can happen is called saddle-node bifurcation, also tangent bifurcation. It looks a little bit peculiar; that is, something really specific. There you have, for instance, one limit cycle and that limit cycle goes to a limit cycle of a totally different shape and form.

You all know this type of bifurcation, every acoustician knows that. If you have a resonance curve and the resonance leans over because of nonlinearity and you have this region of hysteresis, at these two boundaries exactly this saddle-node bifurcation occurs. You come from one side and then this low-amplitude oscillation stops. There is this bifurcation and the oscillation goes to another one -- in this case to this one.

It may not go to another limit cycle, it can go to a chaotic attractor, but then it is not a local bifurcation; then it is a global bifurcation. I will not speak about global bifurcations, because that is a zoo not yet well explored. There you can do a lot of work.

PARTICIPANT: I was just wondering if these bifurcations are generally discontinuous functions of the parameter or continuous functions as you wind the parameter up, say, you might--

DR. LAUTERBORN: In the first case (the Hopf bifurcation) it is continuous, it grows out of a fixed point, and in the second case it is not. There is a jump, really -- it is not really a jump, because if you follow the oscillation in time, it very slowly grows (or decays), but if you go to the steady state, then it is a jump. The motion, critically slows down. It cannot decide whether to stay down there or go up to another amplitude; that is what is continuous somehow. Nature is continuous in that case. In time it is very, very slow. It is almost dead at that point, exactly.

PARTICIPANT: But the attractor itself is --

DR. LAUTERBORN: The attractor itself is totally different. That is discontinuous, yes.

[Transparency 21, upper part]

We have two more bifurcations, and one became very, very famous: the period-doubling bifurcation. It occurs when you have a limit cycle and you get exactly double the period -- exactly. That is a bifurcation and it happens very often, and is very peculiar.

People have observed it all the time, this exact period doubling. For instance, in acoustic cavitation, since the 1950s it has been known that if you turn on the sound field you get -- at a certain amplitude -- exactly $\frac{1}{2}$, whatever you do, to your instrumental limit, and also $\frac{1}{4}$. That was known in the 1950s, early 1950s.

No one had any idea where that came from, no physical idea of what was the mechanism. Also, we at the Institute, we did a lot – Erwin Meier did a lot of cavitation work. We had no idea of what happens, no physical law for that, so we were very eager to study that, to see how it happens.

Any driven oscillator in some part of parameter space has an abundance of period doublings and I will show you them with a bubble later. You can take the laser-rate equations, any nonlinear oscillator you can take, a Duffing oscillator, a van der Pol oscillator, with many oscillators you have an abundance of period doublings, but mostly you can do that only with computers.

[Transparency 21, lower part]

Then there is one very peculiar bifurcation (most people do not know it), a fourth one, called transcritical bifurcation. That is an exchange of stability of two fixed points, a very tricky bifurcation, but there is a very prominent example. If you take the simple laser rate equations, then you have no light -- which is very easy, you turn on and nothing happens -- but then there is a laser threshold and exactly at this laser threshold there is an exchange of stability where you have photons and you have no photons, there are two fixed points for having light and having no light and they exchange stability at the laser threshold. That is a transcritical bifurcation.

[Transparency 22]

These are these four types of local bifurcation and I have produced a few more viewgraphs to show what happens, for instance, at the Hopf bifurcation first. To the left your trajectories spiral in and along that axis c you have what is called your bifurcation parameter. You see the spirals in words going to a fixed point, but when this parameter is changed, at a certain point the fixed point gets unstable and you get spirals outwards to a limit cycle, and the change is continuous. That is what happens in a van der Pol equation.

Then the saddle-node bifurcation. The most prominent example is if you have resonances, nonlinear ones, and they lean over. Then these two points, ω_1 and ω_2 , if you take ω as a bifurcation parameter, these are saddle-node bifurcations, or tangent bifurcations.

[Transparency 23]

Here, just to make the period-doubling bifurcation a little bit more visible, we solved the laser rate equations with pump modulation. Here you have the population difference and here you have the modulation frequency. You know that people want to transmit information on a light carrier and if you do that modulation and you do it in the wrong way you get chaos instead of letting you transmit your signal.

You must not modulate, for instance, too fast. If you have a semiconductor laser, then the relaxation oscillation of these lasers is in a gigahertz range, so if you modulate the laser in the gigahertz range it goes chaotic, so you must stay below that, so that is a certain limit. You cannot go to terahertz frequencies, as people say they want to, with a gigahertz range laser; it will get chaotic. That sets a limit to transmission rate. Most people do not know that and they speak of terahertz rates. You must have very special lasers, of course, to do so, and people learn that from chaos theory. Concerning the transcritical bifurcation; these are calculations from the simple laser rate equations. You have a laser threshold $p=b$. For this $p=b$ there has been awarded a Nobel Prize; they calculated that there should be a threshold, some numbers, but

actually that the pump must be sufficiently strong, must be more than the damping -- we call that damping b and the pump p and normalize. It is just $p=b$ and when you have this, the laser starts to emit light.

If you have the photons, there is an exchange at that point $p=b$ and you get light out. Beyond here, the population difference n is clamped to a static value n_2 . If you are above the laser threshold in this very simple model, the energy from the pump is directly going into photons. It is a very interesting fact.

If you come down with your laser threshold to almost zero, and people try to build that-- you pump your laser and the energy is directly going to photons. That is real, that really happens. These are the four types of local bifurcations and you should keep in mind and a few more global types, this zoo, then.

If you do some measurements and you alter some parameters and you have a nonlinear system you should be aware of these bifurcations; that can happen. You turn on a parameter, you may get an oscillation where, before, there were none. You may get period doubling if you already have an oscillation, -- and that is really peculiar: you can make that experiment and measure that.

You have what we call a passive system that does nothing if you do nothing with the system and then you turn on some mechanism to drive it, somehow to squeeze it periodically, and then it may happen that that system oscillates at a frequency irrational to your driving, so you can drive a system at a frequency totally different from your frequency: two totally different frequencies in an irrational ratio. You can do that, really. I call that an irrational resonance, because you feed energy in -- it is crazy -- at one frequency, and it takes energy out in that irrational way.

[Transparency 24]

Now to another notion that is totally different from all linear things you know, that of the Lyapunov exponents. If you think how is it possible that you have a chaotic attractor, something like that down here that you can measure (the pendulum attractor), how is it possible?

Then you find out something very peculiar: It cannot be, but only in linear terms it cannot be. Again, of course, all these are mathematical terms. Mathematicians found out, especially this person Lyapunov, if you take a volume of initial conditions -- of course, a physicist never would have this idea, we take an initial condition and follow, we cannot take a million initial conditions and follow or we then must make a million experiments -- but mathematicians have developed that idea, that you take a million initial conditions and then you look at what happens.

You follow for a certain time and at an advanced time you look at where does your trajectories stop. The trajectories stop at a lot of points, so that your system is stretched out in one direction, but it is squeezed in another direction and that really happens, as I will show you later in a video.

If that (stretching) goes on and on, then your system will go to infinity and that is what, in linear systems, happens. If a linear system is unstable and you describe it with this exponent, it goes exponentially to infinity, but that cannot happen, of course; it stops before that. What everybody says is, oh, there is some nonlinearity that stops it, but how does it stop it, why a nonlinearity? Nobody tells you. You only read that, that nonlinearity stops it.

In which way does it stop? In a nonlinear system it has a very refined method, it bends it around, this nonlinearity. You expand it and nonlinearity bends it around and in that way it stays finite.

If you go on with this one (the horseshoe $v(t')$) on *Transparency 24*, put this one at the beginning and stretch it out and fold it over, again, you always stay finite, but what you get is all a mess, everything is intertwined, and that is exactly what happens in a nonlinear system, everything gets intertwined.

To show you what that means, if you have two points that are very near by -- you take two points very near by -- those are stretched out and folded over and you find out: If you take two nearby points, they always go off all over. You can take any two points, they separate -- here and here, and here -- you have a dynamical system and those are chaotic dynamics where every two points on your attractor separate.

That is possible. Sometimes you are thinking and thinking and watching and you detect that that is possible. Every two points separate. I will show you. We made a video from that one but it is a summary. I will show you a calculation of a Duffing attractor.

I did the following. I have this attractor transparency 18, $\omega t = 0$, $\omega t = \Delta t$, $\omega t = 2\Delta t$ and so on, until I am one time around $\omega t = 2\pi$, and then it repeats. I plot this attractor a little bit later, so I have 5000 points, these 5000 points and how they evolve in time in different sections, Poincaré sections, all around, from $\omega t = 0$ to $\omega t = 2\pi$.

Then you can observe two points nearby. You look at that and wherever you look -- it is difficult to look at all these points, but wherever you look you see they separate. They always separate.

I will show this video now and Claus-Dieter is operating that. (The video was shown.)

We have a sequence of Poincaré sections of the Duffing attractor calculated, going around, so you see what happens. It is stretched out, you see, turns around, and that goes on and on.

PARTICIPANT: What are the values on the axis?

DR. LAUTERBORN: The values on the axis are x and \dot{x} , so the coordinate and velocity. I call that u_1 and u_2 , that is, the coordinates in the Poincaré sections given by $\omega t = 0$, $\omega t = \Delta t$, etc. That goes on and on forever on a chaotic attractor.

If you go on in time, that happens again and again and again, so such a chaotic attractor must be an incredibly complicated system. If you watch it, as time t goes to infinity, you get one sheet and the next sheet and the next sheet. There is an infinity of sheets, so if you blow that up, you get lines and lines, an infinity of these, so it is an awfully complicated thing already in these Poincaré sections.

Then if you plot all these curves all around, you get something like spaghetti on a dish; you see nothing. That is what happens with a chaotic attractor, and any chaotic attractor must have this structure -- any. In any case it has this stretching and folding property; that is essential.

How to describe that, how to make that, say, mathematically correct? There are several methods. One, if you have everything in the computer and if you measure something, you must make it a little bit different --

DR. ATCHLEY: Excuse me, from what we just saw, you started with a lot of initial states and let it evolve for a while?

DR. LAUTERBORN: Yes, actually what I did is I took one initial state already on the attractor, let it go around 5000 driving cycles and plotted these 5000 points, when $\omega t=0$. These 5000 points is the set of initial conditions. Then I start around with maybe 60 or 100 or so sections for different phases of the driving in between $\omega t=0$ and $\omega t=2\pi$ taking the same data set with stored values from the calculations until one time around 2π and then, to make it not so complicated or having so much calculation, I fitted it to the beginning -- nobody noticed that, but it can go on, it is no problem.

DR. MIGLIORI: Are you going to make any comments about Louville's theorem?

DR. LAUTERBORN: Yes, okay, I have not said anything about that. If you have a dissipative system, then a volume in state space must shrink, if you have a conservative system it does not shrink, and you have no attractors. I forgot to say that. This is valid for dissipative systems only.

There are mathematical theorems that state that properly and how fast it shrinks, you can calculate for these oscillators and so on.

DR. MIGLIORI: Then there is always this problem of how big, or what the infinitesimal volume of phase space you take to measure with.

DR. LAUTERBORN: Yes, but these Lyapunov exponents describe it to some extent, what is going on there. That is a very nice thing, especially the Lyapunov spectrum, if you calculate all of them, and that can be done. You can calculate these expanding and shrinking properties with this Lyapunov spectrum.

[Transparency 25]

What are you doing? You take a very tiny ball of initial conditions and you have access, then, for this ball. If you have 10 of these axes in different directions, then you look at how that evolves. If you have a nonlinear system, different axes evolve in a different manner, one stretches, the other one shrinks and, also, they rotate -- that is very little explored, the rotation (we did a little bit of work on that, on rotation in state space).

Then you can make a definition and how do you do that? You take the radius here, say, this one ($r_1(0)$), and you compare it with this one ($r_1(t)$). You take the logarithm of $r_1(t)$ ($r_1(0)$). The logarithm is very good, if you have the number 1 and you take the logarithm, it is zero. If the number is larger or lower, then you can discern it.

Then, of course, if you go on, t being very large, you must normalize that, you normalize with $1/t$, and that is really all; you compare these two radii. To make it mathematically correct, you have to make t going to infinity and you must go with the radius to zero. That is very peculiar, at first you will not believe that. You take something like a linearization, you linearize, and you work with a linearization, but you take this linearization all along the nonlinear trajectory.

At every point you look at how the linearization looks, whether there is expansion and contraction and so on, and you average over that. That gives you a number for every one of these radii and these numbers are called Lyapunov exponents.

You see, you have as many as the dimensions of your state space, so if it is 10-dimensional, you have 10 Lyapunov exponents and they make up the Lyapunov spectrum. Usually you order them largest one, second-largest one, and so forth, and that is important. You have m Lyapunov exponents and you can calculate them.

[Transparency 26]

Now what can you say about your system if you calculate the Lyapunov exponents? If you have ordered them -- that is the largest one, the largest one is number 1, -- if it is below 0, then you have a fixed point. You calculate, then you know to have a fixed point Lyapunov exponent. If number 1 is 0 and the second one is lower than 0 -- it is very difficult to calculate, that something is really 0 -- then you have a limit cycle. If you take two points on your limit cycle, they go round and round and if you take the average, they stay there, there is no expansion and no contraction and that means exactly 0, and then you have a limit cycle.

If two of these, λ_1 and the next one, λ_2 , are 0, then you have a torus -- very difficult to verify -- and the next one must be below 0.

If the largest one is greater than 0, then you have a chaotic attractor and now we can fix it. We see it, in any experiment you see it immediately, but we can attach a number to that now, how chaotic it is. That is, if λ_1 is large, above 0, then it is very chaotic; if it is not so, it is less chaotic.

The normal way to express that is that if you have a chaotic attractor, then you have what is called sensitive dependence on initial conditions: If you have two arbitrarily nearby initial conditions, then the system, nevertheless, blows up. That is made quantitative with the notion of Lyapunov exponents.

You see, if you have a chaotic attractor you cannot predict anything in the long run. After some time everything is mixed up, you do not know where you are on the attractor, this point or that point. If you have explored the whole chaotic attractor after some time, you can be here, here, or here, the weather can be stormy or sunny, or anything. You cannot predict for a longer time.

I think I will stop here and we have these 10 minutes of recreation.

[Transparency 27]

DR. LAUTERBORN: Now we come to nonlinear time series analysis. If you do some measurement, you normally do not have the state space; that is what is missing. How do we combine theory with experiment? How to get the state space, that was a big problem at the beginning of all this chaos theory. If you have some measurement, how will that fit with all the theory? That is a problem.

[Transparency 28]

If you do a measurement, that is an experiment, you measure -- today you have everything digitized, but that does not matter -- you measure one, usually one single variable or function of a variable (that also is possible) as a function of time, or maybe two or three and then it is already multi-spectral, but usually you have one.

If you are in theory, then you have a trajectory and a high-dimensional state space and does that fit? How can you fit a 1-dimensional signal with a 10-dimensional state space? That was a big, big problem. You can even argue that if I measure only one coordinate, is that enough to characterize that nonlinear system? Aren't we upset and all we do is simple nonsense, somehow, we measure just one variable and then we argue about what may happen, and so on?

The mathematicians worked hard and they came up with a very simple procedure that everybody can immediately do and implement to solve that problem. That is very fortunate and that gave chaos theory, a nonlinear dynamics theory a big input.

[Transparency 29]

That runs this way, very simple. You have a time series of sample data and I call that S_1 , S_2 , S_3 , and so on, you can sample that at a kilohertz rate, a megahertz rate, gigahertz rate, it does not matter. To do the embedding, you take these samples in, say, bunches of 3, 4, 5, you can choose. If you want to have that in 10 dimensions, you can have that, and it goes as simply as with three.

You take the first three, S_1 , S_2 , S_3 , make a vector of that and that is a point in 3-dimensional state space. Then you take the next ones, S_2 , S_3 , S_4 , make a point out of it. You can also take S_1 , S_2 , S_3 , S_4 , and you have a point in a 4-dimensional state space, that is all. You just group them that way.

You also can take S_1 , S_3 , S_5 and put those. Nobody knows beforehand, if you do measurements, how to do it, but you just take consecutive ones, because if you have measured S_1 and S_3 at a lower pace, then you do not have S_2 , for instance, so you have no other values, then, to embed as S_1 , S_3 , S_5 . You get a sequence of these points.

If you do that, and these are real experiments now -- yes?

MR. SMITH: I just want to make certain that that embedding uses consecutive data points, right?

DR. LAUTERBORN: Yes, you use consecutive ones, not mixed, but you can sample at 20 kHz, you can sample at a megahertz; you do not know beforehand, but you take normally consecutive ones.

If you have measured at a megahertz rate and you find out that it is an acoustic system going at 10 kHz, you leave out the other ones; you take every 20 or so, because the other ones are of no use. I will show you why.

We have a pendulum and we do these measurements. We, for instance, measure the angle - you can do that very easily, you have something attached to the axis and you measure the angle. Then you embed the ample data in a 3-dimensional state space; that means you measure at a certain pace and embed that.

It may happen that you get a number of points grouped like that (a) and a number of points grouped like that (b), et cetera (e). In this case (a), you have sampled very fast. Then what happens is S_1 , S_2 , S_3 are almost the same, the pendulum is going at 1 Hz, you sample at a megahertz, S_1 , S_2 , S_3 are the same and, if they are the same, you are along the diagonal.

If you are too fast and you embed and you get just one diagonal, then you are too fast. If you do it approximately right, then everything is, as people say, "unfolded," it blows up, that trajectory (b) blows up, and then you are approximately right.

If you have something going at, say, 10 kHz and you sample it at the rate of 1 Hz, then it usually looks like that (e), all a mess somehow, scattered points. After some probing you find out what your system does. If it is a laser, maybe you must sample very fast. If it is an acoustical system, you need not sample so fast. You must try that. That (b) is approximately

right and these are experimental data and you can play with that if you have a definite system to test your ideas.

We have a video of how the embedding works in real time. We installed an experiment so that we can connect it to a computer and then turn it on, altering a parameter, and the computer does it all, so we immediately have this state space. We also plotted the spectrum to see what we have, what lines in the spectrum we have, and I will show you how that can be done. This was also diploma work at our institute. (The video was shown.)

We had an electrical system with one nonlinear element -- to the right, that one with the nonlinear element. The other things were only capacitors and coils. That is called a Chua oscillator. The nonlinear element has a broken linear voltage-current characteristic.

What have we done? We sampled some voltage at some point of the electric circuit and then we embedded that in three dimensions -- you see three axes here. We plotted the spectrum and that is all online. You can turn the parameter and watch what happens. That was on a Silicon Graphics machine, so you see that period doubling, you have observed that, and now the person turned around -- his name was Merkwirth -- and you see period doubling (you now have period 4).

You also observe that in the spectrum, because if you have period doubling -- somehow somebody disturbed that tape and that goes on to chaos -- so you can have that online if your computer is fast enough. If you have an acoustic system, you can also do that online. You just put in your data and you embed it, you plot it, and you watch what happens, whether there is a period-doubling bifurcation or whether there is a Hopf bifurcation, a, e, something starts to oscillate from a fixed point; you can observe that immediately.

So we did that, for instance.

[Transparency 30]

If you have done that and you have reconstructed your state space, so now you are suddenly from one dimension into five dimensions, what do you do, then? If you are in five dimensions, you have no idea what your attractor looks like, so you must do something about it. The computer calculates how it looks; you need some characterization of your attractor.

Here I have put it down. If you have embedded your data and you see a structure, then it is worthwhile to look into that. How can you see a structure in ten dimensions, for instance? You cannot see it; you must let the computer look through your data to decide whether there is a structure or not. In order to do that, you must do some tasks. You must characterize the structure and that is something static; you have your point set. You do not look into the dynamics; you have your point set. Yes?

MR. HICKEY: How do you choose the dimensions?

DR. LAUTERBORN: You have to make trials, you make 3, 4, 5, 6, 7 -- later on I will give an example. If you do not know your system, there is no way, you must try.

MR. HICKEY: Looking at the time series, there is nothing there to give you a hint to go--

DR. LAUTERBORN: It really looks like noise, your time series, and then you make Fourier, it also looks like noise, but when you do this type of analysis, suddenly it does not look like noise and then you can tell it comes from a deterministic equation. That is the main aim: what you want to know, is this noise or is it some sort of strange attractor and with what properties?

You must do some characterization and one of these is dimension estimation. Even if you embed in ten dimensions and if you have this object, almost one-dimensional, regardless of whether you embed it on a plane, as here, or you embed it in a three-dimensional space, it is always one-dimensional. You must tell the computer how to tell you that this is 1-dimensional, and I will tell you how you can do that.

That is only static, you look at only these points and you characterize, so that is something like chaos static, but you want to have chaos dynamic, so what do you do? You calculate the Lyapunov spectrum from your embedding. I will show you how you can do that.

The final aim, of course, is, if it is not noise, you want to know the law. People are already on the way to having the computer tell you the law. In our times, if you would measure how the planets go around the sun, from this analysis we can immediately write down the law; the computer would write down the law. With such a simple law it writes it down immediately.

That is the ultimate aim and we are already a long way ahead on that; the computer gives you the law.

[Transparency 31]

The first part is to determine the dimension of your object. That runs as follows. I give you, from the many definitions of how to determine the dimension, a very simple one that easily can be understood. You have your point set M . Then you take one point out, point P . You plot a ball around that -- on this plane it is a circle -- then you count the number of points inside. The number of points inside is $N(r)$, because it depends on the radius.

This is $N(r)$, that is, the number of points inside the circle of all these points. Then you shrink your r , shrink, shrink, shrink down, and then you get a function, $N(r)$, and then you try to fit something like r^D . If that fits nicely, then you say the exponent D is the dimension. It looks strange at first, but if you test that on a normal situation -- an example, if you have a line and you take a point on that line, you count the number of points and you go down in r , you see if you have half the radius, you have half the number of points, if you have a quarter of the radius, you have a quarter of the number of points: it goes as the exponent 1.

The computer counts the numbers of points and says, well, that is a dimension-1 object. You can do that, really. If you have two dimensions, you will see that works, also. You count the number of points, you know, if it comes to a smaller and smaller radius, even if it is curved, if you come to infinitesimally small radii, you see it goes as r^2 , so your computer tells you that is a plane, as dimension two.

Now what happens if you have a point set that is a chaotic attractor? That also works, but you come up with a fractal dimension. You find out the object you have has not dimension one, not dimension two, not dimension three, but something in-between, and that works.

You can characterize, with this type of counting of your points, the dimension; you can calculate that. The most famous set is the Cantor set where this has been done, but I will not go into that deeper. You only have to note that this way of counting points gives you the dimension of an object.

If you have done that, you know the dimension of the object. If you have embedded into five dimensions and you find out your object is only one-dimensional, then you have done too much. You could embed that in a lower dimension to find that out.

What is the implication if you have found that out? The implication is that you now know something about your system. For instance, if you have really white noise, then everything should be independent, there should be no law for how to come from one point to the next, no law at all. If you embed that and you embed that in higher and higher dimensions, the computer should never find a structure. It should always come up, if you embed it in ten dimensions, your points give ten dimensions. If you embed it in eleven dimensions, you will find it is eleven-dimensional, etc. Then you know it is white noise.

If it does not, there is some law behind this, some correlation between the points, some law to come from one point to the other one. You know something from that dimension. If you come up with three dimensions, or four, then you know it should be possible to describe your system with a system of differential equations of just four variables; you do not know which ones or which combination of pressure-temperature say, but you know four would be sufficient, so you know a lot about your system from that. You cannot learn that from Fourier transform or from something else; there you learn nothing about your system, but if you do this analysis, you know that, for instance, so you can discern chaos from noise; you have a whole spectrum, 2-, 3-, 4-, 5-, and so on, dimensional.

MR. HICKEY: In your example you had a fractal dimension of 2.35. Would you embed that in a three-dimensional system?

DR. LAUTERBORN: Yes, but it depends on the system. Usually you need a few more dimensions. Mathematicians have shown that in this case $2 \times 3 + 1 = 7$ surely will make it. If you have a nonlinear oscillator, say, like the Duffing oscillator that I have shown, after the embedding in three to four dimensions you come up with something like 2.35, or something like that. You can really measure that at the system. If you alter parameters and you get a different chaotic attractor, you may get a different fractal dimension.

Of course, if you have 2.35, then you must say you need three variables to describe it and it really works, at least in the computer. If you make a Fourier analysis it is just noise.

[Transparency 32]

The next step is to learn something about the dynamics -- this was just statics, statics of these points. How can I get at least the largest Lyapunov exponent from my measurements? Again, you embed your data and you get, for instance, this line in state space from consecutive points of embedding; you now take the consecutive ones.

Then what you do is you look for a nearby point. You have your measurements and then you look, after the embedding, is there a nearby one, so you look through all your database and look for a nearby one. If you have found that, then, of course, you have some topology in to decide what is the distance.

This point ($L(t_1)$) goes that way, this point ($L(t_1)$), in this case, goes to that point ($L'(t_2)$). What do you have to do? According to the definition of the Lyapunov exponents you do exactly what the definition is; you take that distance and compare it with that distance, and you have an expansion then. Or, if you have a contraction, then you also compare this.

You take this length (that means this distance) and divide it by that distance. But after some time, because of this folding over, that will not work, so after some time, a little tiny time,

you set it back. You look again for a nearby point at that time, and that may be that one, and that goes like that to that point.

You again divide the two distances and then you sum it all up all along the trajectory, as many points as you have. If you do this, divide it by the time you had here -- of course, it takes a little while, that was the definition -- then you get the maximum Lyapunov exponent.

Why the maximum? If there are other ones that are not so expanding, you always will go with the largest expansion. If you think about that, if you take these points at random, there is always a part in the direction where you have maximum expansion. That has to do with these subspaces. You can calculate the largest one.

I had this little list.

[return to Transparency 26]

It is very important whether this Lyapunov exponent is larger than zero, lower than zero, and you can calculate that. Then you can decide, even in five, six, seven dimensions, where you have no idea of how to see that, the computer can tell you that you have a chaotic attractor or that you do not have a chaotic attractor.

You can do this type of analysis and that is totally different from everything you had before, all the theory that appeared. Now, if you measure something and until today you have done only Fourier analysis, what do you do? I made a little viewgraph of what to do.

[Transparency 33]

You have your experiment. You measure a time series; everybody has that in the computer. Then you can make a linear time series analysis and, of course, you should do that first, because if it is a linear system, everything works out fine, it is okay.

It may be that you find out it is somewhat noisy, what you measure. In former times you said, well, there is too much noise in my system, I will start all over, again. Today, before you do that, you do this, because, maybe, your system went chaotic and then you have to determine the properties of that system first.

You reconstruct your state space and you try to embed in two, three, four, five, six, maybe going up to ten, because there are a lot of obstacles in going to high embeddings, you need a lot of numbers that you have measured, and so on. When you have done that, you try to make a characterization, dimension, Lyapunov exponents, the largest one; you can also characterize the whole Lyapunov spectrum, that is possible. You can calculate what are called entropies -- there are many more. I have just taken the basic two. You can do all kinds of statistics and so on.

There is also something like nonlinear noise reduction. In your high-dimension space you can do something with your data, totally different things. This here is very dangerous. If you have a time series coming from a chaotic system, you put a linear filter in-between and then embed. That is very dangerous. You can do it, but you may totally destroy your data by linear filtering. There are examples. You have a chaotic attractor, make a linear filter and everything is gone.

This is an interesting thing, your surrogate data. As a preliminary test before you embed your data, you can do the following. You take your data, you make a Fourier transform, then you shuffle all the phases and then you transform back and then you compare these two data sets embedded, and if there is no difference in the embedding, for instance, in the fractal dimension,

and you have done this shuffling of statistics, stochastic shuffling of data, no difference, that is the first indication that your data are simply noise, so do not go through all that, then, because you do not need to.

There are a lot of things to do and, of course, this one is the most important. You want to have some prediction of how your system behaves in the future. A lot of work is going on in that area of how you can make the optimal prediction, the most optimal you can do from your data. People are working strongly on that and, of course, on applications. In the States there is a company called the Prediction Company, where people try to make money from that.

These are the basics of nonlinear time series analysis and we are working on a computer program so that all of that will be automatic. You have your data, put them in, as in a Fourier analyzer, to have a dimensional analyzer, Lyapunov spectrum analyzer, so just put your data in and the computer tries them all out.

DR. MAYNARD: Suppose in an experiment you have more than one transducer and measure more than one time series simultaneously?

DR. LAUTERBORN: People are working on that to put that together. Usually you do not earn much if you make two, because if they are somehow dependent, then you earn nothing. We are working on that with a company in Germany to supervise large power plants, to learn about the minimum number of detectors you need, because they have maybe 100 detectors, and we looked at that and said, "Good Lord, at least 50 are too much," and that has to do with this type of analysis.

You see, you measure just one variable and you go into a high-dimensional space and that works, so you do not need the other ones. It is very peculiar. We measured just one, so if you measure the second one, maybe you gain nothing or maybe you gain something. We are working on that. That is a very interesting question. For instance, in application, how many monitoring devices do you need to supervise something? That depends on the complexity of that system; that has to do with the complexity.

[Transparency 34]

Now we come to some sort of applications. I start out with the theory of nonlinear bubble dynamics. I told you that if you have a nonlinear oscillator and you drive it, peculiar things can happen.

[Transparency 35]

The bubble is a very nonlinear oscillator. This is the bubble and some parameters you need, the radius of the bubble, R , the radius of the bubble at rest, R_0 . There is an outer pressure, P_e , and there is an inner pressure P_i , the density of the liquid, ρ , viscosity μ . I call this ratio χ the specific heat. That is the minimum you need.

You see there are a lot of parameters you have. There is also vapor pressure, there may be some gas and so on. That is a very complex system and a very nonlinear one, a very nonlinear one, because you can, from this equilibrium, compress the bubble to only radius zero but you can expand it 10 to 100 times. That is a nonlinear oscillation surely, if you have just one radius compression in 100 radii expansions. It is a strongly nonlinear oscillator.

[Transparency 36]

Anthony has given you some equations. The Rayleigh equations are the simplest ones. We took a little bit more where the liquid is compressible. This is the Gilmore model. We took this type of equation to calculate how the bubble would behave, bubbles of different sizes at different driving frequencies and at different driving pressures. You need that, for instance, for sonoluminescent studies. If you have a bubble of 8 μm or 5 μm and you have a pressure of 1.3 bar at some frequency, what is the oscillation of the bubble? You have some theory of that and how would that fit with the oscillation observed?

Just a few words, what is extra here in these equations -- "extra" is the van der Waals hard core -- what that means is you cannot compress a bubble to radius zero. Before that you have all the molecules of an uncondensable gas, so you have an excluded volume. That makes a difference.

[Transparency 37]

Now, how to plot the results? Of course, if you take a bubble and you have an initial condition of radius at some time and a velocity or so that is very complicated to compare with other bubbles. You know, if you have a harmonic oscillator and you plot the amplitude response curve -- I did it (actually Dr. Nettin did it all) in the same way but, instead of taking the driving pressure or the driving frequency, which in most cases of application in sonoluminescence is fixed -- because you have a fixed container and you have this one fixed frequency of the basic radiation -- but what can alter (by gas diffusion) is the radius of the bubble at rest, so I plotted that along the axis, the radius of the bubble at rest, R_0 .

I plotted here something equivalent to an amplitude. As this amplitude I took the maximum radius minus the radius at rest and divided that by the radius at rest. That is a very interesting quantity, but something like the amplitude of the oscillation.

What happens? If you take 20 kHz and you see the lowest curve down there, that is for 10 kPa, that is, 0.1 bar (or atm), 10 kPa, 30, 50, 70, and you see the bubble is really strongly nonlinear. At 10 kPa you see what happens, the resonance curve goes over, so these are steady-state solutions.

That is called the main resonance and there is one more, a tiny one that is a resonance that is called a higher harmonic resonance, and there are numbers. You see 1/1, 2/1, and so on. One over one, that means you drive the system with a frequency and the system answers with the same frequency, and that is the main resonance here. What is that, 2/1 harmonic resonance? That is really the resonance, the system, all these resonances with 1, 1, 1 in the denominator answer with the same frequency as the driving.

This 2 in the nominator means that there are 2 oscillations of the bubble in one driving period, or 3 or 4, up to six here that fit, and every time they fit there is a resonance, so there is an infinity of harmonic resonances.

Here is a resonance, 1/2. What happens there? You have 2 driving periods and just one oscillation and then it repeats. The first subharmonic resonance, it is called. This happens in this way in all nonlinear oscillators; we found that everywhere. If you have a nonlinear oscillator, you have all these resonances and all these rational numbers and so on. This can be made much more formal mathematically via so-called torsion numbers and we worked on these.

Three over two, what is that? They are resonances in-between with different numbers of drivings. In this case you have two periods of the driving and your bubble makes three oscillations that also gives a resonance, so you can get, at sufficiently low damping, resonances for every rational number.

If you have a linear oscillator, it starts here, goes up, you have one main resonance, you have only $1/1$, nothing else, and it goes down here. You can normalize that; you have only one curve normalized. In this case you cannot normalize because you see this lower one and the upper one are totally different. You cannot normalize.

[Transparency 38]

Here are a few curves, just to give you an impression. That is the driving, two periods, and that is from the main resonance. That is what it looks like. You have something like a large excursion phase, a collapse, large excursion again. That is the main resonance, a typical picture for the main resonance. If you measure something like that, you are in the main resonance.

[Transparency 39]

If you are in the other ones, as I told you, it looks like that, what I have here. These are solutions from $2/1$ and $3/1$. So these are 2 periods of the driving, and two oscillations per period of the driving, and that repeats. If you are in the $3/1$ resonance it goes like that, and you see the higher resonances have a lot of oscillations in-between in one period of the driving.

[Transparency 40]

Here is one solution from the $3/2$ resonance. You can make a whole plot from every resonance of what that looks like.

[Transparency 41]

I will go very quickly through that. If you go higher up with your driving, what happens? I call those response curves, no longer resonance curves, because you encounter things other than just resonances. Again, we have radius here and the maximum oscillation amplitude here. In this case, the calculation was done in the following way. You see there is a curve here (solid line) and, simultaneously, a point here (dashed curve). What does it mean?

We did the following. You can calculate this curve in the following way. You start at $10\text{ }\mu\text{m}$ and then you go up and up and you see what happens. You stay on the lower curve, go up here, and you go through all this, and then along there.

If you come from above, from $200\text{ }\mu\text{m}$, you come down, you go this way -- here, here, down here, and so on, but that is not the whole story. That means there are two coexisting attractors -- that may happen. That is where the curves lean over (you have 2 of these jumps, dotted lines).

What is in-between and is that the whole story? Because we have jumps inside also. We go inside and there back and forth. There is a whole mess already and that is at 0.8 kPa . This part has been resolved at a higher resolution

-- down here -- and also checked for more coexisting oscillators.

You see it really starts to become peculiar. We have not yet plotted just simple response curves; that means the maximum radius. We plotted all the coexistent attractors. We have here

a bifurcation diagram. We come along here, for instance, come along here $((2\mu)_1)$ and then it jumps down. We jump to a period-2-oscillation (P2) in this case.

Then that combines here and goes that way. If you come down here $((3\mu)_1)$, you see what happens. We plot for every period of the driving one point, exactly as we did before, now with the chaotic attractor. What we did is: this is also a maximum amplitude but it is the relative maximum amplitude in consecutive periods of the driving, a very complicated plot.

If that does not repeat, you get a lot of these points and you see that that is chaotic (7/2). You see that at 0.8 kPa a normal bubble gets chaotic, it shows period doubling -- this is a period doubling -- period doubling to chaotic oscillation, jumps in chaotic oscillation, very, very complicated. We have done many more calculations to resolve that and it is just extremely complicated.

[Transparency 42]

A few curves from coexisting solutions. Depending on your initial conditions, you can have that (upper diagram) or that (lower diagram) steady-state oscillation, and there may be many, many coexistent attractors.

[Transparency 43]

This is one of these chaotic solutions. It seems quite regular but it is chaotic.

[Transparency 44]

Now we go ahead and we are heading for the small bubbles at high driving for these sonoluminescence parameters. We go up in pressure and here we have 90 kPa and we go from 30 to 7 μm and you see what happens. Here we have resonance number 20 and if you go up and up in pressure here you see these marks (lower diagram) -- I will show curves of what that looks like -- you go up 100, 105, 110 kPa, you see what happens. These chaotic regions come down to smaller radii and expand; the chaotic solutions expand almost all over the parameter space of the radius of the bubble at rest.

[Transparency 45]

These are curves, what they look like from the 20th resonance. If you count that properly, 1, 2, 3, 4, 5, you go all around, there are 20 maxima, 20 of these oscillations. Here, also, 20 of these. You can tell from these oscillations in which resonance you are.

[Transparency 46]

If I go ahead with these "20"-resonances and go up in driving, you see what happens -- all have 20 of these after bounces or rebound oscillations.

[Transparency 47]

Now I go ahead with driving. Beyond about 1 bar, where you tear the liquid, you see here, again, is the 20/1 and a very peculiar thing is happening if you go up. Suddenly you get very large relative expansion rates. Before, you have seen, you come down here like that, but now you go up and come down. You go up and go down. That is very peculiar.

You have your main resonance and everything was going down but suddenly, at about 120 kPa, you get something like an extra response of the bubble. You see how that goes up. This (the radius scale R) is a logarithmic scale. A logarithmic scale is very good if you have period doubling and all that; then it shows up beautifully.

If you plot it on a linear scale, here you see what happens.

[Transparency 48]

I have done it here for 130 kPa and 140 kPa and all over the range from almost 0 to 100 μm . That is the main resonance (1/1). That is how the bubble behaves for small bubbles, it is the same driving. It is a very peculiar behavior. What happens is all this chaotic stuff ceases here. If you go higher in pressure, it retreats, even.

[Transparency 49]

I can show you chaos net cost in this view graph for 150k Pa. We call that a giant response. That is the main resonance (to the left in the upper diagram) and you get a real giant response for small bubble -- resonance is not good, it is a not resonance; resonance is all these tiny things here in-between -- but you get a giant response and that is what is called the inertial range of response of your bubble.

If you drive it beyond about 1 atm or a little bit more because of surface tension, then nothing hinders the bubble in its expansion; you just tear it apart and the expansion is dominated by inertia. That is how it comes about.

Yes, the chaos, the chaos all retreats, the bubble gets stable, again.

[Transparency 50]

Here I have these plots in one diagram. Here is the driving and the driving goes up from 0.7 bar to 1.7 bar. For a bubble of 7 μm the calculations were done. What is very peculiar: if you calculate the curves, you see that the afterbounces are a long burst of oscillations. They are getting smaller with pressure increase, also in theory, not only in experiments but also in theory, the afterbounces get smaller, and only afterwards, for very high pressures -- 1.7 e.g. is a very high pressure -- they come up again. That has to do with the damping by acoustic radiation and is in the equations.

[Transparency 51]

One last picture of these calculations and then we again make a little pause. What happens if we do it all over, again, not for 20 kHz but for 30 and 40 and 80k Hz, all I did before?

DR. MATULA: If you started out with that bubble for this curve, 10 μ , on the far right-hand side, and you are driving the system at the same driving pressure, 150 kPa, are you saying the bubble is going to dissolve and, in response, going to get bigger?

DR. LAUTERBORN: Yes, if you would do that experiment, and at 20 kHz your 8- or 5- μm bubble has a larger amplitude-- you see what happens here. If you would go up with your frequency and it would dissolve, yes, I would say, because the amplitude of oscillation gets lower.

DR. MATULA: So it would follow that 20 kHz line, the response line, and actually --

DR. LAUTERBORN: It should, up or down, depending on the gas content. but it settles down at a point according to your gas content: here, here, on here on the right shoulder of the giant response, it depends on the gas content and it settles down at lower radius and higher oscillation amplitude at a lower gas content, it gets smaller in radius with a lower gas content and acquires a higher amplitude. That higher amplitude is the trick to avoid complete dissolution.

In these calculations diffusion is not in, because if you have diffusion in, it is totally different. If you start here, here, or here, you end up at the same radius, all these bubbles

dissolve down to the same radius. If diffusion is in and the bubble loses or gains gas, it will settle down at a certain radius and that would be the solution including gas diffusion. No curve of the kind shown one stable, only points on them depending on gas content.

We have done these calculations but another German group, Lohse, has also done this calculation. If you include the diffusion equation, then you can calculate more precisely what happens.

Here I wanted to show the dependence on frequency. If you go up in frequency, you see that all the chaotic stuff comes down to smaller radii. That is very peculiar. It clears off if you go up in driving. We should do some more calculations on the relation between driving frequency and how far these chaotic solutions reach and what you can do. That should be, I think, important for some sonoluminescence experiments, what happens if you go to different frequencies.

I will stop here and then the last part will be our experiments. Thank you.

[Transparency 52]

DR. LAUTERBORN: Now we come to the experiments, the best part. All the other parts were preliminaries so that you understand what we are doing.

[Transparency 53]

Everybody likes this, so I will go ahead with this one. You can put a bubble in the interior of a flask and then drive it, say, at 25 kHz with power amplifier and impedance matching network -- that (impedance matching) is fine. Then here (at the flask) you drive it at some fundamental resonance and then you measure on the oscilloscope some output or put in a hydrophone.

[Transparency 54]

We heard on Monday morning from Anthony what happens is this bubble starts to luminesce -- there it is. You know the story, Crum invented that cell (we call that a Crum cell in our laboratory).

[Transparency 55]

You have seen that it also happens in a Coca-Cola glass but it also happens in rectangular spaces and that is especially bright, we found out, really nice bright luminescence.

[Transparency 56]

Also cylindrical -- we have done it in different types of forms -- and in this case with two bubbles. Also we can place 1, 2, and also 3 to investigate more on that, just to show that.

I have a little video on what that looks like if you turn on the sound field, put in a bubble and then observe what happens with a CCD camera. C.D.O he and put it on videotape. (The video was shown.)

(Something on video before) This was just normal water poured in, a bubble, and then look at that. The bubble is dancing around there and making light. That is what normally happens.

You can also see that with the naked eye. You just look at that and it immediately starts. If you lower the gas content, then the bubble may stay very stable at one point. I say nothing new here, but we have a nice picture of one of these luminescing bubbles.

[Transparency 57]

We have an exposure time of 13 hours, so we set up the experiment in the evening and then went home, opened the shutter of a camera and the next morning looked at the result and the camera collected all the light and it was totally overexposed.

DR. MATULA: What are the dimensions?

DR. LAUTERBORN: The dimensions from here to here (height of the picture) are about a millimeter. The bubble was not totally stable but it was just standing on a table, not suspended, nothing of that sort, just standing there, and maybe cars are going around (we are very near a street, unfortunately), so the bubble moved a little bit according to what happens, but it is stable over a long, long time and makes a lot of light so that we can photograph it. I think that looks pretty nice.

[Transparency 58]

Several people have measured the radius time curve of such a bubble. We also did that. To photograph it, I think Bob Apfel at Yale was the first to do it that way. We are a little late now in this business but, nevertheless, we photographed the bubble, but in this case we did not do this high-speed photographing during one cycle but from one and the next and the next, and so on, cycle with a shift in phase versus the driving -- you photograph it every shifted time.

You see what happens. The bubble goes up in size slowly, very slowly, and then it collapses; it really goes down and collapses. Then there are a lot of afterbounces and then it repeats. That is one full cycle around.

I plotted down our data from the experiment, 21.4 kHz, 1.32 bar, and the size of these frames and we took every 500 nsec -- this was a shift between the frames -- every 500 nsec we took a picture. In that case (the collapse) you can even place some pictures in-between, but this is enough, I would say, to take radius time curves from this sequence.

[Transparency 59]

If you do that and compare it with calculations, we came up with this one diagram here. These circles are all experiments from the frames. The solid line is the calculated curve. There was some fitting. At 21.4 kHz we took 1.32 kPa; we took a static pressure; the bubble radius at rest is a little bit difficult, a tiny bit of fit; surface tension, as water should have (I am not sure whether R. Geisler remeasured that). He measured viscosity, because there was a little bit of glycerine in, so it is 0.0018 Newton sec/m²; density of liquid like that.

The biggest fit we did was the polytropic exponent. We took 1.2, because nobody knows what is inside the bubble or whether you even could take a simple polytropic exponent. Then we get this fit along here. You see, from all the afterbounces we can almost guess in what resonance we are, but we are surely in this giant response.

You can count these oscillations up to 12 if you go along here; it is a little bit of a mess. What we guess is, you are always around 20, maybe, in that resonance. Also, that (some oscillating excursion is the after-bounces) is only a tiny thing up and you hardly can notice that in the curves. You can discern all the resonances from these oscillations.

[Transparency 60]

Additionally, we measured the shock wave coming out and this is all from R. Geisler, a diploma student of mine, he did that all. I think he put all that on the Internet and you can take it from there, all the sequences; you can make a video from that.

Here you see the shock wave coming out from here to here to here. The shock waves are much faster than the bubble oscillation, so we increased the speed. We took an every 10-nsec shift to make a photo. For the same bubble as before we have this shock wave going out and you see the shock wave gets dimmer and dimmer, and that is okay. It expands, so it gets dimmer and dimmer.

That is the shock wave outside. You cannot look inside. Inside that is really, really tiny, it is just a point on a photo and, if you want to look inside, what happens is you see the shock wave that is going out. You want to look inside, but the light is deflected; you cannot look inside. You cannot even properly determine the radius because the light is deflected from the compressed water. That is a very difficult problem.

People have done that also with light scattering -- Putterman's group did that -- and then you have the scattered light and you can transform it into a bubble radius. You come up with totally similar curves.

That was single-bubble dynamics, and we also have done similar things, and they are really similar things, with laser bubbles.

[Transparency 61]

You can have the same with laser-produced bubbles. You focus laser light into the liquid, then the bubble expands, collapses, and that looks totally similar, except that the bubbles are bigger. Again, you can see bubbles growing large and then suddenly collapse and make some afterbounces, so that is the laser produced bubble in water.

We have done a lot of this work, so that is companion work to all the sonoluminescence and, also, we measured the light for these bubbles just to show that very similar things can be done with another type of producing bubbles.

These were single bubbles. Much more work has been done by us on many bubbles: what is called acoustic cavitation.

[Transparency 62]

I told you that was the problem we started from. We produced acoustic cavitation by making a strong sound field in a liquid and measured the sound output from the liquid. In the other picture I showed you *[Transparency 13]*, there you could see all these branches and streamers of bubbles, and I have a video to show that. (The video was shown.)

This is a dynamic process and if I show you just one picture, you will see nothing moving. There is a lot of dirt inside -- this is just tap water, so all these tiny bubbles. Radiation pressure puts them into the pressure anti-node. All these bubbles are collected and if the bubble is large, it goes off to the pressure node.

PARTICIPANT: Can you use this to de-gas?

DR. LAUTERBORN: That can be used to de-gas, yes. To a certain amount you can de-gas that way, yes.

What happens here, you go up in pressure amplitude and it is plotted there. Now it is high pressure, beyond 1 atm, 1 bar, 100 kPa, and you see what happens.

DR. GARRETT: Is that self-luminescent or is that being illuminated?

DR. LAUTERBORN: That is illuminated, but if you photograph that without light, it looks very similar. We have done that also in self-luminescing light.

You see what happens if you go up and up with your driving. There is a whole sequence of events that start if you go up with your pressure amplitude. One special thing is you have all these branches of bubbles slowly moving.

DR. HARGROVE: Are those bubbles or are those the trajectories of bubbles, these lines going out?

DR. LAUTERBORN: These lines going out, quite clearly these are bubbles coming in.

DR. HARGROVE: But it is the trace of a bubble moving below your time resolution?

DR. LAUTERBORN: In this case I would say yes, because it is photographed at video speed, so that is a streak line. If you photograph that with a few nanosecond illumination, you just have spots there. If you do that constantly, you see the spot going, and they are going mostly from outside to inside.

In the last part, if you have high pressure they go outside, inside -- inside, outside, they turn around; they do not want to go to the middle. We think that has to do with the bubble oscillation being strongly nonlinear and not being attracted to the pressure anti-node. They are repelled in a strongly nonlinear case. We had a paper on that, that there should be a threshold for nonlinear oscillations, that they go to the pressure anti-node. If you are higher in pressure, they go off.

[Transparency 63]

This is a very old viewgraph, because there is a transient recorder. Now you do that directly in the computer but in former times, a few years ago, you needed a transient recorder because the memory cards were too slow in the computer. Today you go directly into the computer.

You drive the system with some frequency and you have a hydrophone and you store what you hear, what the transducer hears, in a computer. What have you then measured?

[Transparency 64]

I will show you that. Also, I have a video for that. That is what happens. If you start at a low driving, you put in a single frequency and maybe there is some distortion in your system; then you have higher harmonics. That is the spectrum of what you put in, your frequency f . I put a 1 here and that is the actual frequencies. We had around 20 kHz, for instance.

This amplitude is the amplitude of the spectrum. If you go up from here to here in the driving amplitude, then all of a sudden exactly $\frac{1}{2}$ appears. You see there is no noise below. If you go up, $\frac{1}{4}$ appears and all of its harmonics. After some time you get some sort of broadband spectrum.

Because we had a period-doubling sequence, we have shown that we end up with chaos. I will show you a video where we took our data and we go up with the driving amplitude and made a sequence of Fourier spectra for different sound pressure amplitudes. (The video was shown.)

This is very old, 1981, you see. That was on film and we replotted it on videotape, so a little bit got lost that way, but you can see it nicely. Here is the time. Here is the excitation level

in volts. We simply go up with the volts to the cylinder. That is how time passes by, plotted there.

Today you do this easily with the computer but in 1981 you had an 8-bit computer, so it was very difficult. You see here we are in $\frac{1}{2}$ harmonics, $\frac{1}{4}$ now. You see some modulation products and you see it is not totally stable, but $\frac{1}{2}$ after some time is always in, and that was what people noticed very early. If you turn on a sound field in a liquid and you go to high amplitudes, you always have $\frac{1}{2}$ of your driving.

That actually was the starting point of all of our investigations. We wanted to know what was going on there.

MR. WAKELAND: Why did you go through that sweep so quickly?

DR. LAUTERBORN: That is to have the same physical system. If you are going up too slowly -- we did it in $\frac{1}{8}$ of a second. We found that was, for this case, about the optimum. If you go up too slowly, then your system heats up, your bubble system is moving -- you have seen that moving slowly on a scale of seconds -- so you must go up at a certain speed.

If you do that too quickly, then you go too quickly through the bifurcations; you smear them out. The system has no time to settle to $\frac{1}{2}$ to $\frac{1}{4}$ to $\frac{1}{8}$, so the ideal would be you know how the system behaves, then go somewhat more quickly to $\frac{1}{2}$, but then very slowly to one-eighth, because it needs to settle down to that attractor.

There is a small range where you can do that.

MR. WAKELAND: If you just set it at a particular high driving amplitude and watched it for 10 seconds, you are saying the spectrum would not be -- stable?

DR. LAUTERBORN: It then looks as it did the last time in the movie. You see $\frac{1}{2}$, you see some noise coming up, you see $\frac{1}{4}$, so it fluctuates. That is what happens. You do not see this nice sequence.

After some time we found out how to do the experiments, we computerized everything, and then tested it all. We also worked with pulses rectangular and Gaussian pulses, and -- it is almost, in any case, similar, so we have the best conditions for our system in the case shown that give a clean sequence.

In this case we have up to $f/4$. We have seen it to $f/32$. You can then discern all these lines and very regularly we can have $f/8$. The experiment was first and that was without calling that chaos; it was called acoustic cavitation noise. We had that whole sequence. At that time, in 1981, if you look back at what was known of chaos theory, it was almost nothing.

In 1976 there was the paper of Feigenbaum on his number, but if you looked at that paper you would not have noticed that that had some connection with our work, theoretical work, on maps and on the properties of maps behaving in some limit. After some time it turned out that we had to read that paper for some sort of explanation, because period doubling is something that happens almost everywhere, it is somehow universal, and that is what he actually showed, that in certain types of maps with a quadratic maximum and some extra properties like unimodality -- a map with only one maximum and some properties of its derivatives -- have the same scaling number when you look at this period doubling. I will not go into details.

We learned that something is with this period-doubling sequence and started to read and after some time we came up with these Poincaré sections. With Poincaré sections you can fit

these experiments to maps. If you have such a map, it should behave the same way, because that is universal, and after some time we found out that that is the right theory.

Around the years in the 1980s the embedding came up.

So we said, oh, wow, now we have not only the period doubling but also we can embed the data and find out if it is really chaotic or is it broadband noise.

[Transparency 65]

That was the nicest embedding we have. Holzfuss has done the experiments. He graduated and got his Ph.D. for all this type of work. He graduated with embedding and got his Ph.D. for Lyapunov exponents.

That is a beautiful embedding in a 3-dimensional space and it shows up a clear, very clear, structure. We also embedded, of course, our whole data series so that we could see all the period doubling, because it is period doubling and then this structure. You see 4 bands here.

[Transparency 66]

There was always this question and I can answer that now. When you do your dimensional estimation, what do you do? You have your data set. You choose the dimension you want to embed your data in. You take 2, 3, 4 -- it goes to 10 -- and you take your data, embed them, determine the dimension, as I have told you, count. You put a circle around your sphere and count the numbers when you see it shrinks, and then you calculate.

That is what came up with these last data. You came up with 2.48 and you see -- the variation is not plotted, but it is not that large if you have 2 or 3 one dimension; it is really in-between. You see what happens, you embed your data higher and higher and it levels off, and that is exactly what I told you. If you embed this object here (some pointing stick), 1-dimensional, into a higher and higher space, it stays 1-dimensional, the computer knows it; it shrinks the ball and counts the numbers of your object and tells you: it is 1-dimensional.

These are our data and what are these data? Of course, we needed a test. This was really one of the first experiments where something was done that way, so we needed a test that was sensible. These data are real noise. We took it from some electronic apparatus -- real noise -- and handled the data the same way. We digitized it the same way as our cavitation noise data, the same formalisms, the same embedding, everything the same, same algorithms.

You see what happens with true noise. It goes up here (the noises). There is a little bit going off, it is 9 instead of 10. I had exactly that question in the last break. If you do not have enough data, then you cannot fill the space properly and then if you count the number, you do not have enough and that gives you a lesser number. If you had more data, it would go up.

That works properly and we are pretty sure that this is different. I do not quite remember whether that was (the filled triangle) -- something else. There are tests, some tests.

[Transparency 67]

These were the dimension estimation data. Of course, you can do the following. You go up with your voltage and you always embed and determine your dimension. We have made videos of that, you can do that online directly, and then you find out your data go like that and suddenly they go up.

It is really interesting. Period doubling is in this area down here; it does not show up, and that is quite clear. If you have a line, at that period doubling you have 2 lines, but it is still 1-

dimensional. Only if you come to chaos do you get that jump. That is really an indication that you suddenly have something chaotic -- it jumps.

We also found in other cases that if you plot what we call energy bifurcation diagrams, the energy does not change in a system if a period doubles -- very peculiar. It does not make a jump, it is totally smooth.

[Transparency 68]

Then we calculated the Lyapunov spectra. That also works, the whole spectrum, where one Lyapunov exponent is positive -- not really big, but really sensible. One is about zero. That is okay if you have a continuous dynamical system. One is always zero, and then one is less than zero.

Because we have embedded in three dimensions and that was enough, we have only three, but you see there are some more here. If you embed your system in, say, five dimensions, then you have five Lyapunov exponents, but your system is only three-dimensional, so you have two extra.

That is the sort of problem -- and that is really a problem, because we found out, the extra ones need not be the slowest, there can be a mix, and it is somehow difficult to find out which are the real ones, but that can be handled, sometimes.

[Transparency 69]

This is the last slide I want to show you. All these methods of nonlinear dynamical theory you can apply to real physical systems and if you should encounter where somewhere it looks noisy, please do that and perhaps you will detect something.

I want to stop here, 3 hours is enough for that, and I thank you.

PARTICIPANT: I have got a question about higher dimensions in Poincaré sections. My understanding is that normally the Poincaré section is used just to reduce the dimension by one, so that the map is one dimension lower than the state space you are working. Would anything stop you from successively applying Poincaré sections to maps one dimension higher to reduce it to something you can actually draw on the paper or --

DR. LAUTERBORN: Yes, but there is a problem. You do not know how your trajectory goes to phase space, so you do not know where to place your section. In this driven system, it is very easy. You take your phase ωt somewhere. Hopefully you can take the whole plane at different times, so you have what is called a global Poincaré map.

In most cases you have only a local one and then you must shift properly that local one, and in higher spaces that is all a mess because you do not know how your trajectories go or whether they will ever come back or circle around somewhere and then you are in that place and find nothing.

That is a very, very difficult question and only in very specific cases can you do more than one.

QUANTUM MECHANICS MINI-TUTORIAL

Julian D. Maynard
Department of Physics
Pennsylvania State University

DR. MAYNARD: Thanks. Doing a 4-semester course in an hour and a half is really impossible, so I am going to talk fast and I probably will not finish, but I spent a fair amount of time working on these lecture notes so, hopefully, they are pretty self-contained. If I talk too fast and do not finish, please read the lecture notes, because I think they are pretty easy to read and they have most of the concepts in quantum mechanics, and there are a lot of concepts and aspects that you need to learn about. You have to do it sort of all or nothing, so let's get busy.

[Transparency 1]

First of all, the purpose: Why am I giving this lecture? If you are learning acoustics, you might wonder why you need to learn quantum mechanics. There are a number of reasons. They involve the same wave phenomena. In fact, acoustics, optics, and quantum mechanics are all just different versions of what I call wave mechanics. What you learn in one field you can apply equally well in other fields.

Things that I learned in quantum mechanics that I apply in acoustics include group theory and the effects of symmetry on wave functions; evanescent waves; dispersion; group velocity and the spreading of wave packets; the effects of degeneracy; and perturbation theory (and there are probably some other things). Those things you can learn in acoustics or optics but you may not, but you certainly learn those in quantum mechanics.

Another application is if you are studying the attenuation of sound, as in air, that involves molecular collisions that have to be treated with quantum mechanics. This is a large sub-area of physical acoustics.

Solid crystals have quantized sound waves that are called phonons. If you want to understand solids, like the heat capacity of solids, you have to understand phonons. Also, phonons in a solid interact with electrons and the two together can show effects of pressure, magnetic fields, electromagnetic radiation, may do Raman or Brillouin scattering.

People in acoustics study these effects. You can study high-temperature superconductors and other new exotic materials, quasi-crystals. You can also do surface acoustic waves for thin films, and so on. All these things, in order to understand them, you need to understand quantum mechanics.

I mentioned superconductors, but you can also do acoustics in other macroscopic quantum systems, such as superfluids. In fact, you can have a pretty successful career studying the acoustics of superfluids -- for example, Isadore Rudnick.

[Transparency 2]

Let's start off. Usually courses in quantum mechanics take a look at some history, so we will take a look at some events that led to quantum mechanics. First of all, in 1901, there was Plank's theory of blackbody radiation. A blackbody is a solid that can be in thermodynamic equilibrium with an electromagnetic radiation field.

If you take a look at the amount of electromagnetic radiation as a function of its wavelength, then you get a distribution that is given by these points here. Classically, people tried to understand this distribution. They had a classical theory that could fit this side, but then it would blow up that way. They had a theory that could fit this side, but then it blew up the other way. There was no classical theory that could fit this entire distribution curve.

What Plank did, he assumed that solids had elastic modes that could have only discrete values of energy. You know an elastic solid can vibrate at discrete frequencies but you can have any amplitude and any amount of energy in that mode. Plank said no, you can have only certain values of the energy. He said that the amount of energy was given by a constant times the frequency of the mode. If you think about it, that is a very radical idea. The constant that was involved is now called Plank's constant (symbol h) and it is 6.6×10^{-34} joule seconds.

Some textbooks will treat blackbody radiation by saying that light behaves like particles. That is a little misleading. Plank actually did it right. The energy of an elastic solid is quantized in Plank's constant times the frequency.

[Transparency 3]

Another event was the photoelectric effect by 1905. Here you have a metal in a vacuum. You shine ultraviolet on it and electrons are ejected from the metal surface. If you have a battery and an adjustable potential and you measure current, you can determine the energy of the electrons that are photoemitted from the metal surface.

What you find is that the emission of the electrons begins almost immediately, even with a light intensity of only 10^{-10} W/m². In fact, if you do a little calculation, you find that conservation of energy is not even satisfied here.

Also, the energy of the electron is proportional to the light frequency. Einstein had an explanation for this. He said that light waves behave like particles (which are called photons) and the energy of those light particles are given by Plank's constant times the frequency of the light. This was another radical idea. The amazing thing was that the same constant (Plank's constant) appeared as in the blackbody radiation, so this is a pretty amazing result.

Unfortunately, many textbooks say that the photoelectric effect shows that light must act like a particle. This is completely incorrect. You can explain the photoelectric effect by assuming that light acts like a particle and the electron acts like a particle. Einstein was right.

You can also explain it by assuming that light acts like a wave and the electron acts like a wave, so because of that explanation you cannot just assume that the light has to act like a particle. Most of the textbooks are wrong in that respect. Here is an article that gives you a good write-up on the misuse of the notion of a photon.

Although this conclusion is wrong, having light waves act like particles, at least historically, suggested that particles can act like waves, and that is a good conclusion.

[Transparency 4]

Another event was Rutherford scattering, 1911. This did not actually require quantum mechanics itself. What Rutherford scattering showed is that atoms are made up of electrons orbiting a heavy nucleus. You need that result for the next event, which was Bohr's theory of atomic spectra.

If you heat up some atoms and they collide, they will emit light, but the light they emit occurs at only discrete frequencies. Bohr's explanation of the discrete frequencies was that the electrons orbit the nucleus classically but at only certain radii.

He said that the allowed radii are such that the electrons' angle of momentum is an integer multiple of Planck's constant (this time divided by 2π). So again it was another appearance of Planck's constant. This was a pretty exciting time for people with this one particular constant showing up everywhere.

[Transparency 5]

Another event was Compton scattering in 1923. This was x-rays scattering off electrons and the x-ray wavelength has shifted a little bit. Again, you can explain it, assuming that the light or the x-rays behave like particles, but it is not necessary. As with the photoelectric effect, the conclusion that light has to act like a particle is incorrect.

Finally, de Broglie gets it right, that particles behave like waves. Take note, graduate students: This was de Broglie's Ph.D. thesis in 1924. He said that an electron with a momentum p and an energy e is really a wave with a wavelength λ given by Planck's constant divided by the momentum and the frequency ν given by the energy divided by Planck's constant, so Planck's constant appears, again.

This idea explains all the preceding experiments. I showed here an electron and an atom. You can get the discrete frequencies by assuming the electron has to fit an integral number of de Broglie waves in its orbit.

He explained the previous results and he predicted that electrons should diffract like waves. In 1925, shortly after, Schrodinger gets a theoretical basis for particles behaving like waves, the Schrodinger wave equations, and of course I will say more about that.

In 1927 there was Heisenberg's uncertainty principle, and I will say just a little more about that. Finally, in 1927, there is the experiment by Davisson and Germer, where they took at electrons and scattered them off the surface of a crystal and showed that electrons diffract just like waves and produce an interference pattern. That was the clincher for quantum mechanics and particles behaving like waves.

[Transparency 6]

In the history, there were some misinterpretations; for example, that light waves have to behave like particles, which is not true. A lot of the misconceptions and misinterpretations are still being put in textbooks. In this viewgraph, I list some of the myths and the mystique of quantum mechanics.

Quantum mechanics is a new paradigm. Classical physics is just an approximation -- you should remember that. Quantum mechanics is a new paradigm and you should just accept it.

Some of the myths and mystique of quantum mechanics: There is this notion of wave-particle duality, where electrons may behave as either waves or as particles. However, the law of physics is the Schrodinger's equation and it is a wave equation, so everything has to be explained in terms of waves.

The classical particle picture is just an approximation. There is an analog here. In optics, if you have a bunch of lenses, you can calculate how they behave by using ray tracing, but if the rays have to go through a small hole, then you know that you really have to do diffraction and use what is called geometric optics to really explain what is going on. People who do ray tracing

know that it is just an approximation, but they do not go around saying there is a ray geometric duality for light waves. It is really not necessary.

Classical mechanics was around for a hundred years or so and it worked on its level, a macroscopic level; I guess for some reason people are reluctant to give up the notion of point masses behaving classically.

Some old textbooks used to talk about paradoxes that arise from quantum mechanics. You have an electron that goes through two slits, you cover one slit, and funny things happen. There are no paradoxes in quantum mechanics. You get paradoxes only if you insist on giving objects particle-like attributes or if you use detectors that do not obey the laws of physics. Otherwise, there are no paradoxes.

The uncertainty principle: This has significance only if, again, you assume that things behave like classical particles. If you just accept the fact that particles are really waves, then you have no problem and you do not really need an uncertainty principle.

Lately there is something called quantum chaos. This is really different from all this, but the Schrodinger equation is linear, so there really is no chaos. Again, this is just a consequence of an approximation. In fact, it is an old effect that people knew about in acoustics.

[Transparency 7]

Let's get to a valid picture of quantum mechanics. Question?

MR. WAKELAND: When people say that when you do scattering, say, between electrons, that as far as we know the electron is a point particle, no one has found the radius of the electron. You are saying that is just all hooey and no one should even be talking about the radius of anything? Is that what you are saying?

DR. MAYNARD: Sure, yes. (Laughter) We have to solve the Schrodinger equation to understand it correctly, right? So, what does that have to do with --

DR. MIGLIORI: What do you mean an electron is a point particle?

DR. MAYNARD: You are asking him, right?

DR. MIGLIORI: Yes that is nonsense. An electron is what it is. It obeys what he is about to show you.

DR. MAYNARD: Here is the valid picture of quantum mechanics. It has 2 parts. The first part is you solve a boundary-value problem, which is mathematical, rigorous, and has a unique solution. There is no funny business at all. The interesting thing is that after step 1 absolutely nothing happens.

Step 2, to make something happen, you have to make a measurement. This is where it can get a little tricky. First of all, a measurement involves a very large number of complicated elements. You might think that the minimum simplest element is a graduate student with a lab notebook; there are serious theories for measurement that say that human consciousness has to be involved before you have a measurement. That may or may not be valid but it is not really necessary.

A system with many elements is sensitive to small perturbations; there are rigorous mathematical theorems about this. No system is totally isolated, there is no shield for gravity waves, so there are always small perturbations. Roger Penrose has a theory that says that since you cannot shield against gravity waves, you are always involved in measurements, and

everything that happens to us involves gravity waves. The total mass of the universe really helps determine what happens as we proceed with our lives and our measurements.

Because of all this, the best thing you can do is using the results of 1, the boundary-value problem, to calculate the probable outcome of a measurement, because it is always going to be perturbed. The best you can do is calculating some probability. This was put forward by John von Neumann, a very smart guy.

There are other pictures of what measurements in quantum mechanics do. There are parallel universes and all sorts of things that you can use for scripts for Star Trek, but this is quite sensible and works very well as an explanation for quantum mechanics.

There is a writeup in *Physics Today*, October 1991, "The Decoherence in the Transition from Quantum to Classical." It is a very good article.

DR. LAUTERBORN: May I add something? You have not used the word "nonlinearity" for measurements, and I think that is central, because if you, for instance, take ψ squared and the absolute value, then that is a nonlinear operation. I think that is always involved and that makes it interesting. Otherwise, everything would be proportional and linear and nothing would happen. I think the central thing in measurement is a nonlinear operation, always.

DR. ATCHLEY: Could you repeat the comment?

DR. MAYNARD: In quantum mechanics, I will show later, when you calculate probabilities, it involves a nonlinear quantity, ψ squared, and because of this nonlinearity, again, complicated things happen and you are left with just being able to calculate the probabilities of things.

To do quantum mechanics, you have to first understand classical mechanics, so I am starting at the beginning here and we will go through this fast.

[Transparency 8]

Newton's law: $F=ma$ or dp/dt , where p is linear momentum, mass times velocity. The important thing in quantum mechanics is angular momentum, $r \times p$ -- and I should have written here that Newton's law takes the form, the time derivative of angular momentum is torque, which is $r \times \text{force } F$ (we will need that later on).

Kinetic energy is $1/2mv^2$. Potential energy is a line integral of $F \cdot dr$ and, at least on a microscopic level, all forces are conservative. Conservation of energy, kinetic plus potential, is a constant if you take all the forces into account.

[Transparency 9]

An important part of classical mechanics involves electromagnetic radiation (at least for quantum mechanics this is very important). A force on a charged q is given by this expression, where it involves an electric field, E , and the velocity of the particle crossed with a magnetic field, B . You can get the electric and magnetic fields from potentials. There is a scalar potential, ϕ , and a vector potential A , and the fields are given by those expressions.

Here is something you may not have seen in a mechanics course. You can establish an electromagnetic force potential. This potential is $q\phi q/c(A \cdot v)$, and the force from this potential, the x component here, is $-dU/dx$, just like an ordinary potential. Then there is a term, the time derivative of the derivative of the potential with respect to v_x , so this now is a velocity-dependent potential. That may be something new to you.

[Transparency 10]

In advanced classical mechanics courses you are introduced to generalized coordinates: r here is the coordinates of one particle and it is given as a function of these generalized coordinates, q_1 through q_N , and may be a function of time.

You may have M particles, so for, say, the N th particle you have this expression in terms of generalized coordinates. These are three-dimensional vectors; so you have $3M$ coordinates in real space for the particles, but because of constraints in a mechanical system, you can have the number of generalized coordinates may be less than $3M$.

One example is suppose you have a frictionless hoop with a bead on it -- like this -- but this hoop is rotating around at some angular frequency ω . You can describe the position of the bead just with some angle, θ . With it going around on the hoop, it is constrained to be on the hoop and the hoop is going around, motion is quite complicated, but there is only one generalized variable, and that is this angle θ .

We can generalize to have velocity-dependent potentials and there could be explicit time dependence there. Now, for doing these complicated problems you can use Newton's law, but it gets pretty complicated. There is a systematic way of doing it using the Lagrangian.

The Lagrangian is the kinetic energy minus this potential energy. The Lagrangian is a function of generalized coordinates and their time derivatives, using the notation that means a time derivative. There is the Lagrangian and then, equivalent to Newton's law, are the Lagrangian equations. It is the time derivative of partial L with respect to \dot{q}_i , partial L with respect to q_i minus that equals 0, and for each generalized coordinate you will have an equation, so you get a whole set of equations to solve.

MR. WAKELAND: So are there N particles or not?

DR. MAYNARD: There are M particles but only N generalized coordinates, because they are equations of constraint. These equations are second order in time.

[Transparency 11]

With the Lagrangian you can define generalized momenta: p_j is partial L /partial \dot{q}_j . An example of a generalized momentum that is very important is if you have that electromagnetic potential. Then p_x is m times \dot{x} , which is the usual thing, but then you have this term, q/c times the component of the vector potential.

Then you define the Hamiltonian. The Hamiltonian is a function of generalized coordinates and generalized momenta (and possibly time) and it is a Legendre transformation, it is the sum over j of $\dot{q}_j p_j - L$, and you have to use this to eliminate the \dot{q}_j 's in favor of the p_j 's.

An example of a Hamiltonian, if you have an electromagnetic field: Ordinarily you would have $1/2 m \dot{x}^2$, but $m \dot{x}^2$ now is going to be $p - q/c$ vector potential squared plus q times the scalar potential. So, this is important in quantum mechanics.

With the Hamiltonian, Newton's law becomes \dot{q}_i is partial H /partial p_i and \dot{p}_i is minus partial H /partial q_i . Again, these are equivalent to Lagrange's equations and Newton's law, but the reason why theoreticians like this is because you have twice as many equations but they are only first order. If you look at numerical recipes, if you have a set of second-order equations you want to solve, step one is to convert them to twice as many first-order equations.

[Transparency 12]

Some conservation laws in classical mechanics: Suppose, perhaps because of symmetry, the Hamiltonian does not depend on one of these generalized coordinates, so p_j , which is minus partial H /partial q_j , if H does not depend on q_j , that is going to be 0, so p_j is a constant in time.

In particular, if L does not depend explicitly on time, then the Hamiltonian is a constant. Under these conditions, where, say, the transformations to the generalized coordinates do not depend on time and you do not have velocity dependence in the potentials, then it turns out that the Hamiltonian is $T+V$, or the total energy (remember, the Lagrangian was $T-V$). The Legendre transformation under the right conditions gives you total energy.

MR. POESE: Just for perspective, the whole point of the Hamiltonian and the Lagrangian are to get these generalized coordinates under control?

DR. MAYNARD: It is a systematic way of solving mechanics problems. It is nice, you can just blindly turn the crank and get the answer.

DR. MIGLIORI: The Lagrangian has a nice interpretation. If you are rolling a ball down a hilly track, the speed of the ball at any particular time is determined by L , because as it comes up the hill, $T-V$ is smaller as it goes down. That is the basis for the WKB approximation later.

PARTICIPANT: Is that like a particle?

DR. MIGLIORI: Yes. It works the same way, so that you can figure out its absolute speed by just looking at L .

DR. MAYNARD: Okay that is it for what we need from classical mechanics. Now we have to take a look at the notion of vector spaces and function spaces.

[Transparency 13]

This could be a whole math course in itself. It is what I found as one of the most difficult parts of quantum mechanics, dealing with vector spaces and operators, and so on. It is easy to write down, but then you are left stranded as to exactly what it is you are supposed to calculate.

You are familiar with just Cartesian vector space. A point in 3-dimensional space can be represented by a vector, and the vector is given by the sum of components times some unit vectors. You have an inner product, usually known as the dot product, an inner product between two vectors, $A \cdot B$, which is the sum over the product of their components.

You can also vector spaces that are function spaces, where the axes and the vectors themselves are functions (just like this function here). You can define an inner product for these function vectors and it is usually an integral of one function complex conjugate times the other. There may be some wading function and then you do an integral.

There is the notion of orthogonality, where the inner product between two vectors is 0. You can have unit vectors, where the inner product is 0, if they are different; if it is the same vector unit it is normalized to 1.

Then there is the notion of operators. An operator acting on some vector in the vector space gives you a new vector. An example for a function space is just taking the derivative with respect to x . For example, if this is the sine of x , then the operator gives you the cosine of x .

You can do a matrix representation for these operators by taking the inner product of one unit vector with T operating on another unit vector. Then there is the notion of the expectation value. That is the inner product between the same vectors of the operator P .

Then there is the notion of a commutator. This is just notation, $[P, Q]$. That is defined to be, first, you apply the operator Q and then the operator P and you subtract from that what you would get if you did the operator P first and then the operator Q . That is the commutator.

[Transparency 14]

Now we can start going through the postulates for quantum mechanics. Formal theory: At a time t_0 the state of a quantum mechanical system (which we will give the label "A") is given by a point in a linear vector space, which we write ψ_A , and this is at time t_0 . It may be a function, but we are talking about vector spaces now, so I can just leave it as this symbol here. We are picking on time.

Then the state is normalized at all times, so that the inner product with itself is always 1. Then you postulate that there exists a time translation generator, which is a linear operator $T(t, t_0)$, such that the vector at a later time, as given by this time translational operator acting on the initial condition. It is really just a definition.

MR. POESE: What does that mean? Is that the position?

DR. MAYNARD: It is a vector in a vector space, period. You can define an operator $H(t)$ such that the time translational operator at a slightly later time, dt -- if you do a Taylor series expansion, it is $1 - iH(t)dt/\hbar$. Why the $-i$, why the i and why the \hbar ? Right now we do not know what H is, but when we find out what it is or how to get it, we will know why there is an i and an \hbar here. So, this is 20-20 foresight.

If you take a look at just a function of time you can do a Taylor series expansion, $\psi(t+dt)$ is given by this, the time derivative. If I take this and this and plug it into here, then I get that $H(t)$ operating on $\psi(t)$ is $i\hbar d\psi/dt$. This, now, is based on definitions, but that is a big equation in quantum mechanics.

[Transparency 15]

Here is where we really get somewhere. Take a look at the dynamics of the expectation value of an operator P . The time derivative of the expectation value -- I just write out what the expectation value is (here is the time derivative) -- I have a product of two functions of time here, so I use the product rule.

I let the time derivative act on this one plus the time derivative acting on this one. The time derivative from that previous viewgraph is going to be $H/i\hbar$ operating on the wavefunction or operating on the vector, so I will have that one from here.

Over here, when it acts on the one on the left -- remember, there was a complex conjugate, so I pick up a minus sign. If you look at this, this is just $1/i\hbar$ times the expectation value of the commutator of P with H . This is straightforward.

Now we come to the definition or equivalent of Newton's law in quantum mechanics, and it is based on something called the correspondence principle. What we demand is that if you take the limit as \hbar goes to 0 (i/\hbar), the commutator of P with H , without the limit you would get the time derivative of the expectation value. What we are going to force is that if you take the limit as \hbar goes to 0, you get the time derivative of the classical quantity. That is the principle.

Let's apply this requirement by letting the operator be a generalized coordinate. If we do that (and this is just q_i), the classical version and from Hamilton's equations that equals the partial derivative of the classical Hamiltonian with respect to p_i . If we let the operator be p_i and

then if we evaluate this by this principle, this has to be the classical p_i ; but, by Hamilton's equations, that is minus the partial of the classical Hamiltonian with respect to q_i .

All of these conditions can be satisfied if we let H be the classical Hamiltonian, but with the generalized coordinates and generalized momenta replaced with operators that satisfy these commutation relations. You have to think about this a little bit, but if you think about it, if I had this commutation relation and I plug in these operators in the classical Hamiltonian, when I do this, you can see the $i\hbar$ cancels out and when I take the limit I am going to get the classical result.

You have to think about that a bit, but that is the setup, that is how you do quantum mechanics. You take the classical Hamiltonian, replace the coordinates and momenta with operators and you have to find operators that satisfy these commutation relations.

MR. WAKELAND: When you say this is like Newton's law, do you mean just in the sense of being very --

DR. MAYNARD: It is the fundamental. From Newton's law you can derive all of quantum mechanics.

[Transparency 16]

One of the things that was always a little confusing to me in quantum mechanics is that there is more than one way to define operators that satisfy those commutation relations. For example, there is the coordinate representation, where q_i goes to q_i (just multiplication by q_i), p_i goes to $-i\hbar$ times the derivative with respect to q_i .

There is the momentum representation, where p_i goes to p_i (just multiplication), q_i $+i\hbar$ times the partial with respect to p_i . In fact, if you work things out with function spaces, and so on, you will find that these representations are Fourier transforms of one another.

There is the Heisenberg representation, where the q_i and p_i are represented by matrices (remember, in function spaces you can represent operators as matrices).

Then there is also the sort of no representation, where you use formal theory of vector spaces and write everything with symbols to do all of quantum mechanics. The trouble with that is you have some expression and you wonder what do I do now. At least here, you know you take a derivative, which you know how to do.

Let me give you an example of a coordinate representation, where we have x and p_x , and let's take a look at the commutator, which is an operator, operating on some function, $f(x)$. The commutator is going to be x times p_x (which is this thing) $-p_x$, with the x here.

I can take this derivative, and I will get this. Here I have got the derivative but now I have a product of two functions, so I have to use the product rule -- so you get these two terms. You will notice the first term and the last term cancel out and leave you with $i\hbar$ times $f(x)$. You have $f(x)$ here on both sides, so you get that the commutator of x and P_x is $i\hbar$, which is the commutation relation we needed. So, it works, or you can make it work.

In 3 dimensions with these sorts of definitions you find something that appears in the classical Hamiltonian (remember, kinetic energy is $p^2/2m$) and that is going to become $-\hbar^2/2m$ times the Laplacian, just using something like this definition.

[Transparency 17]

Suppose we have the classical conditions where the Hamiltonian is the total energy, kinetic energy plus potential energy. That is $p^2/2m$ plus some potential, which may depend on position

and time. Then if we use the coordinate representation, we get that H acting on a state vector, $\psi(r,t)$ as $-\hbar^2/2m$ Laplacian acting on this function plus the potential times the function (this has only coordinates in it) and that equals $i\hbar d\psi/dt$. This, then, is the Schrodinger equation.

Now, suppose the potential does not have any explicit time dependence, it depends on coordinates. In that case, you can use separation of variables so that the state function is just $\psi(r)$ times $e(-i\omega t)$. You plug that in and you get this equation, which is the time-independent Schrodinger equation, $-\hbar^2/2m$ Laplacian, and that is equal to E times $\psi(r)$, where E is $\hbar\omega$.

There is really no difference between E and ω , just two different symbols, tradition has two different symbols being involved here, so this is a differential equation and there can be some boundary conditions.

For a closed system, E will be a discrete eigenvalue, or quantized energy level. It is going to turn out that the expectation value of energy is going to turn out to be this E anyway, so it is a quantized energy level. The eigenfunction $\psi(r)$ is called a stationary state for the time-independent Schrodinger equation. The eigenfunction for the lowest E is called the ground state. Discrete eigenvalues are countable, so they can be labeled with integers or sets of integers and those integers are called quantum numbers.

[Transparency 18]

Let's do some examples.

MR. WAKELAND: When you just said all those things about quantized energy levels and such, that all depends on -- that does not come out of this time-independent Schrodinger equation by itself?

DR. MAYNARD: Right.

DR. WAKELAND: I mean, they could be not quantized?

DR. MAYNARD: That is right. In fact, I guess for the free particle here they will not be.

One-dimensional free particle: There the potential is just 0. It does not depend on time, so you can write the time-independent Schrodinger equation this way. You can multiply through by $\hbar^2/2m$ and you get the Helmholtz equation, where this k is $1/\hbar^2$ of $2mE$. Classically, momentum is the squared of $2mE$, so from this expression you get the momentum is \hbar times k . You can also get that with the operator P .

Solutions to the Helmholtz equation, solutions for the wavefield, are some complex amplitude times $Ae^{i(kx-\omega t)}$. These are waves. The waves have a wavelength λ , $2\pi/k$, which is \hbar/p -- Plank's constant over p . The frequency ν , which is $\omega/2\pi$, is the energy eigenvalue E divided by Plank's constant. Those are the de Broglie conditions for particles acting like waves, so we have theoretically derived de Broglie's hypothesis.

ω here, just by definition of the symbols, is E/\hbar from this expression, $\hbar/2m$ times k^2 . In fact, from the governing equations, when you get this solution for the free particle, as a relationship between k and ω , so ω is some function of k , and that is called dispersion. For common acoustic waves we find from the governing equations, that ω is usually the speed of the wave times k . That is a special case of linear dispersion.

[return to Transparency 17]

In classical acoustics, you have a Laplacian and then you will have a second-order time derivative. In quantum mechanics, you have the Laplacian and a first-order time derivative.

[return Transparency 18]

As a consequence of that, where, in classical acoustics, you have ω , you have linear dispersion, in quantum mechanics we have here that ω goes like k^2 , so it is nonlinear (or it is not linear) dispersion. Dispersion is a fundamental difference between free-particle Schrodinger waves and common acoustic waves.

You might be able to rig up an acoustics system that has this dispersion; I have not really thought of one.

DR. MIGLIORI: Jay, you certainly could rig up a real system that has the acoustic dispersion for electrons, for example, electrons in bismuth have ω equals $c_0 k$.

DR. MAYNARD: Right, this was just for the free particles.

DR. LAUTERBORN: May I add something? You have here a quantum mechanic wave equation and the acoustic wave equation, and perhaps one should say where the difference is. The difference is that in quantum mechanics you have this i , this imaginary unit. That really makes a difference. That makes it first order in time; nevertheless, you have wave solutions.

DR. MAYNARD: Yes that is what I was pointing out here.

DR. LAUTERBORN: Yes, and that also makes the dispersion relation different and, therefore, you have difficulty to see it directly in acoustic waves, the same dispersion relation.

[Transparency 19]

DR. MAYNARD: Let's talk about wave velocities and dispersion. If you have a function like this, $Ae^{i(kx - \omega t)}$, basically trig functions, the argument of the trig functions is ϕ , which is called the phase, so let's address the question of what is the relationship between x and t such that the phase is constant.

If the phase is constant, that means the derivative is 0, so if we take the time derivative of this, we get the k times $dx/dt - \omega$. You can solve for dx/dt . This is the velocity. This gives you the velocity of points of constant phase, and it is given by ω/k , which is v_p , the phase velocity, so points of constant phase move with velocity ω/k , the phase velocity.

Now suppose we have a wavefield that is given by a super position of two waves, cosine of $kx - \omega t$ and cosine $k + \Delta k$ times x minus $\omega + \Delta \omega$ times t . I can use a trig identity and write it this way. Now it is a product of two trig functions. This k bar and ω bar are pretty much the same as k and ω for small Δ s and that is a carrier wave -- this thing here that is oscillating rapidly.

This term here is an envelope; it varies slowly. If you look at the form of this Δk times x minus $\Delta \omega$ times t , that is moving with a velocity $\Delta \omega / \Delta k$. If you take the limit, that is $d\omega/dk$, the derivative, and if you have dispersion ω is a function of k , and that is v_g , the group velocity. This envelope moves with the group velocity. In ordinary acoustics, we have linear dispersion, so the group velocity equals the phase velocity. In acoustics, you can maybe miss worrying about group velocity.

For more general cases when you have dispersion the group velocity is different from the phase velocity.

[Transparency 20]

Once you get these solutions, or eigenfunctions, the general solution is a super position of those. The general solution may be an integral with some complex amplitudes $A(k)$ and you are integrating across over all possible k . In this case, for an infinite system it is not discrete, so we

integrate over k . I wrote here $kx - \omega(k)t$ to remind you that there could be dispersion or could be dispersion that is not linear.

Let's suppose that $A(k)$ is peaked around some k_0 , as shown here. That is essentially the carrier spatial frequency. If I take the wavefield here and evaluate it at $T=0$, then this goes away. If you look at this, this is the spatial Fourier transform, or it is a Fourier transform of $A(k)$.

You should know from signal processing, if I take the Fourier transform of this, I will get something in real space that looks like that. It is a wave packet. This is the shape at $t=0$, so it has some envelope like this. That is the initial condition.

[Transparency 21]

Just some comments: From a theorem for Fourier transforms, the product of the width of the curve in k -space times the width of the envelope in x -space is a constant. That is essentially the Heisenberg uncertainty principle, but it is a rigorous theorem for waves. It is really nothing new; if you accept the fact that particles are really waves, then there is nothing new here.

If the curve in k -space is a Gaussian, it goes like this, and the envelope in x -space is also a Gaussian. Having Gaussians minimizes this product of the widths. Also, you can do calculations with Gaussians, so you find Gaussian wave packets a lot in quantum mechanics. Remember that $A(k)$ is peaked at the value k_0 .

MR. HARDIMAN: What is the significance of the product of the width system analyzed?

DR. MAYNARD: If you talk about the uncertainty principle it gives you the best possible measurement of, say, position and time, or position and momentum, I guess.

DR. MAYNARD: Let's define κ like this, then we can take this initial condition that we had and write it in terms of this κ . You want to take note of that form.

But now I want to go back to the original expression for times other than 0 and what we want to do for this $A(k)$ peaked at k_0 is a Taylor series expansion of the dispersion around k_0 (this ω_{k_0} is just ω evaluated at k_0). The Taylor series expansion has $d\omega/dk$ times κ , but $d\omega/dk$ is the group velocity.

Here is the exact expression for the wavefield at later times but now it is written with the κ . Now we are going to take this $\omega(k)$ and use the Taylor series expansion. When you do that, you get this expression. There is a phase factor here, but all this stuff right here is just the initial condition, but evaluated at this argument, because there is a phase factor and initial condition with that argument.

If you take the modulus of this, if you calculate any expectation values, the phase factor goes away and you are left with the original envelope evaluated with this argument or, in other words, the envelope moves with the group velocity.

MR. POESE: What is the significance of the k_0 ?

DR. MAYNARD: It is just the carrier wave for a wave packet.

MR. POESE: That is the wave number of the carrier wave for the wave packet?

DR. MAYNARD: Yes. It is a carrier spatial frequency -- that is this stuff here -- then there is the envelope.

This says the initial condition moves unchanged with the group velocity, but this was only an approximation. You have to remember that with nonlinear dispersion wave numbers that are higher or lower than k_0 move with different velocities and that causes the wave packet to spread.

[Transparency 22]

Here we go. It is time for some demonstrations and you should know how difficult it is to get demonstrations to work, so let's see what happens. [Computer demonstration]

There are two wave packets. Actually, they are kind of cut in half, because they are $x=0$. The lower one has no dispersion; the upper one has dispersion. [Draws on transparency]

Here is the dispersion as a plot, so here is ω as a function of k . Linear dispersion, a straight line. Then let's take some point here to be k_0 . When I put in dispersion, what I am going to do is let the dispersion go like this. The carrier frequency has the same phase velocity for the pulse without dispersion and the pulse with dispersion. At the spatial frequency, it has the same phase velocity, the same value of ω/k .

MR. WAKELAND: You have your straight line there. Are we calling that dispersion or not?

DR. MAYNARD: It is dispersion. $\Omega(k)$ is called dispersion. In this case, it just turns out to be the special case of linear dispersion -- that is the definition I got out of the book.

I used to think dispersion was when it was not linear but apparently, that is not right (at least in the book I read). [Computer demonstration continues]

The bottom one has no dispersion; the upper one has dispersion.

MR. POESE: By no dispersion, do you mean linear dispersion?

DR. MAYNARD: I am sorry, yes, linear dispersion.

The upper one has some nonlinear, or not linear is probably better.

So, we turn time on. You can see, in the top one -- remember, it curved up, so higher spatial frequencies travel a little faster. You can see here that there are higher spatial frequencies in the front here, because they are going faster, so we are left with lower spatial frequencies here.

This one, with linear dispersion, is traveling unchanged -- what we usually learn in ordinary acoustics, that pulses travel unchanged on a string.

Now you should be able to see that lower frequencies are slower, higher frequencies are faster. That is the spreading of the wave packet that you hear about in quantum mechanics.

I rigged this thing up so that the center frequency should have traveled at the same phase velocity, yet this one got ahead. Why is that? Here is your first homework question. Why did that one get ahead?

PARTICIPANT: It travels at the group velocity.

DR. MAYNARD: Right, it travels at the group velocity and the linear curve went like this. The one with not linear dispersion went through like this. Even though they had the same phase velocity, the group velocity is the slope. The linear one has a slope like this and the other one has a higher slope at that point, so it goes faster.

MR. POESE: How does this relate, again, to quantum mechanics?

DR. MAYNARD: In quantum mechanics, you often represent the old classical point masses as wave packets to get them from A to B. This is what a particle is in quantum mechanics.

DR. FORTUNKO: If you were to take the upper one and determine what the peak would be in terms of the time, can you do that? Could you recover the dispersion from that?

DR. MAYNARD: Yes, I guess so. I would have to think about that a bit.

DR. WAKELAND: So if you have the bottom one, there is just one velocity there, I mean all the velocities are the same, so it is very clear which velocity it is. In the other one you have

all these different velocities, so somehow you are saying there is one particular one that is the velocity at which the packet moves, right?

DR. MAYNARD: Let's go back to the other slide.

[return to Transparency 19]

The phase velocity is ω/k and that is the same at this point where they cross. The group velocity is the derivative, so the derivative of this one is like that. The derivative of the one with the not linear dispersion is like that. The group velocity is higher, so the envelope travels faster.

MR. WAKELAND: My question is, imagine that you had a distribution curve that had a peak but was not symmetrical. To determine what k_0 was, would you take the highest point or the mean point or what?

DR. MAYNARD: However, you want to define it. It is arbitrary. The distribution curve is the thing that is really important. That is why the uncertainty principle really does not matter, because you are dealing with waves, and if it is a distribution function, then it is a distribution function. Just accept it.

[Transparency 23]

Some more examples: There is the piecewise constant potential field. Here the potential is V_i , a constant, or just a number when x is between x_i and x_{i+1} . For that piecewise section the Schrodinger equation looks like this, with a V_i here. You can rewrite that with a k_i (k_i is $1/h$ bar times the square root of $2m(E-V_i)$). In each section where the potential is a constant, you have a wave number of k_i like this.

MR. POESE: Are these quantized or not or is this like the free particle?

DR. MAYNARD: We have not said what the whole potential looks like yet. This is just a general method for treating piecewise constant potentials. We need a particular example to see if things are quantized or not.

For a second-order equation like this with the k_i you have two linearly independent solutions and each one has a complex coefficient like this.

In some regions, E may be less than V_i , so this is a negative number; we are taking the square root of a negative number, so k is imaginary. In those cases the wavefield will look like $e(-kappa x)$ and these are called evanescent waves. You can also have $e(+kappa x)$, which is just an evanescent wave where the k is in the negative x direction. These are not effervescent waves; that is what you get on beer.

The coefficients, A_i and B_i , are found by satisfying conditions that $\psi(x)$ and its derivative $d\psi/dx$ will be continuous at the boundaries between these piecewise constant sections.

Satisfying all the boundary conditions may result in solutions existing for only discrete values of E . This is what we were just talking about and those become the quantized energy levels.

[Transparency 24]

Here is an example. Here you have potential of 0 to minus infinity. At $x=0$ it jumps up to a value of E_0 out to plus infinity.

Then let's take the case where E is bigger than V_0 . There is an implied boundary condition at minus infinity and plus infinity. On this side we have an incoming wave and a reflected wave

and over here we have only a transmitted wave. We do not have any wave coming back from plus infinity.

The solutions, then, are incoming wave e^{ikx} , reflected wave e^{-ikx} , with a complex coefficient \tilde{R} ; k -- over here -- is given by $1/\hbar$ bar $2mE$. On this side over here you have the transmitted wave, either the e^{iqx} with complex coefficient \tilde{T} -- and q is $1/\hbar$ bar $2m(E-V_0)$. We are assuming that $E-V_0$ is positive in this case.

If you equate ψ and $d\psi/dx$ at $x=0$ from these two solutions, then you get two equations and you use those two equations to solve for \tilde{R} and \tilde{T} , which are the complex reflection and transmission coefficients.

These are not the energy transmission and reflection coefficients (that is a little different), but if you have two equations and two unknowns you solve for \tilde{R} and \tilde{T} , then you are finished. The results are that you have waves on both sides. [Computer demonstration]

[Transparency 25]

Here is the step. The energy I still let be where the center of the wave is, which is above the step. As time goes on -- this is going up and down because, remember, there is an incident and a reflected wave, and you get partial standing waves in that case. That is what is making this go up and down like this.

At some point it looks like the transmitted wave is bigger than the incident wave, which may sound wrong, but the actual energy transmission coefficient has a factor of q^2/k^2 , which in this case is less than 1. So, the actual energy transmission coefficient will be less than 1.

DR. FORTUNKO: So the transmitted wave almost holds to the constant amplitude, but the reflected wave does not?

DR. MAYNARD: Right.

[Transparency 26]

Remember, this was the case where E was bigger than V_0 . If we do the k^2 , E as less than V_0 , then on the transmitted side we get the evanescent waves. [Computer demonstration]

Now you can see that the energy is a little below V_0 on this side over here, so we have the incoming wave, but then notice that it is matching the wave value and the derivative of the wave. Over here, we have the exponentially decreasing solution, the evanescent wave. As time goes on, it does that.

Notice we pretty much have a standing wave over here and the reflection coefficient -- if you look the complex \tilde{R} and \tilde{T} are complicated now, complicated complex functions, because q went to iq . But if you calculate the energy transmission coefficient, it is 0. It has the real part of q in it; q is now imaginary. The energy reflection coefficient is 1.

The evanescent waves -- there is a wavefield over here, the wavefield is not 0, but it is not propagating any energy out to infinity, so all the energy has to be reflected.

DR. FORTUNKO: The solution on the right-hand side was not evanescent? In addition to the amplitudes on the left-hand side moving up and down, were the zero crossings also moving?

DR. MAYNARD: Yes, there is a partial wave; it has partial nodes. It has anti-nodes but only partial nodes, so the zero crossing is actually hunting back and forth, absolutely right.

DR. MIGLIORI: The energy density on the left is constant, even though it does not look like it.

DR. MAYNARD: That is right, if you time average.

DR. MIGLIORI: Well, you are showing only one component.

DR. MAYNARD: That is right. I am showing only the real part here.

DR. MIGLIORI: If you did the whole thing, you would find that the modulus is constant on the left.

DR. MAYNARD: That is interesting. As an acoustician, I tend to like plotting the real part. We are going to get to this in a minute.

[Transparency 27]

The next example for the piecewise constant potential is where you have 0 potential over here, then it goes up to a value of V_0 for distance a , then back down to 0. If we have the total energy E or the parameter E up here, then we have propagating waves here, propagating waves here, and propagating waves here (they just have different wave numbers, they are different wavelengths), and not too much interesting happens.

The interesting case, of course, is when E is down here, so you have propagating waves here and here but evanescent waves in this barrier. On the left, you have incident and reflected, with this k . In the middle you have waves going back and forth, where q is given by this, and q is now going to be i kappa, if E is less than V_0 . On this side, we have transmitted propagating waves with k , again.

There are lots of textbooks that will calculate R tilde and also T tilde for you, but if you want to do a computer demonstration, you have to calculate E and F -- not a fun project -- and they are complex. You have to get amplitudes and phases. It is a big mess, but it can be done.

The solutions: You will have a propagating wave over here. The value and derivatives match, but you get evanescent waves decreasing here, but they do not decrease to 0, so you get some transmitter wave here.

PARTICIPANT: Can they go to 0 inside?

DR. MAYNARD: If the width is finite, then it will not go all the way to 0, but if this is anywhere bigger than 1 or 2 wavelengths, it might as well have gone to 0, you cannot measure it. It is exponential decay. If k is like $2\pi/\lambda$, if the distance is 2λ , that is $E(-4\pi)$, which is 10^{-6} ; at 2 wavelengths it is down by 6 orders of magnitude.

MR. POESE: These are all for infinite lengths, right, so far?

DR. MAYNARD: So far, but we are going to do a wave packet right now.

[Transparencies 28-34]

This is what the single k_0 solution would look like. [Computer demonstration]

Here comes the wave packet, for that case with the barrier.

Here is the wave packet, so it has a number of spatial frequencies, but I have set the barrier height so that the barrier height, really, is 1% higher than the highest spatial frequency, so all the wave components in the wave packet, all the spatial frequency components, are going to be evanescent in the barrier. The barrier width is $3/4$ of a wavelength.

Here comes time. You will notice that at any instant I am matching wave and derivative to get the evanescent wave, but through a finite barrier it does not decay to 0. It wiggles on the other side of the barrier and there allows there to be a wave on the other side of the barrier.

Here the wave packet is mostly reflected, but because of the finite width of the barrier, some wiggles got there, and in quantum mechanics, this is called tunneling through the barrier. You can do this with classical waves. You can have a wave incident where you would have total

internal reflection but there is an evanescent wave and, if you have a thin film, the evanescent wave can get through and you will get some real propagating waves on the other side.

MR. POESE: So in effect, this electron has now split. Some of the energy is over here on the right-hand side of the barrier and some is on the left-hand side of the barrier?

DR. MAYNARD: No. What happens is, if I did a measurement, there is some probability the measurement will show that it is over here, there is some, much smaller, probability that the electron will be over here if I do a measurement. That is what it means.

MR. POESE: But only at one of those two places?

DR. MAYNARD: It is just what I said. The wavefunction looks like this and the measurement has some probability of showing that it was here and some probability that it was here. That is all I can say about it. There is no point electron that I can say is here or there. I can talk only about the result of a measurement.

MS. ZOU: It is like 1 over there and 99 on the left.

PARTICIPANT: I think there is some confusion. The interpretation is the probability of finding a point electron and that is what is diffracting? This seems to be sort of the standard interpretation.

DR. MAYNARD: In a minute I am going to do the simple harmonic oscillator and then show more about what a measurement is going to do. Right now I am doing just part one, which is finding the solution to the boundary-value problem, so we will get to part two in just a minute.

DR. MIGLIORI: The wave packet is -- let's see, the high end of it has frequency components that dribble off exponentially.

DR. MAYNARD: Yes, I have a finite cutoff.

DR. MIGLIORI: Yes, you made a finite cutoff and --

DR. MAYNARD: And I made my barrier higher than that finite cutoff, so they are all evanescent in this example.

MR. WAKELAND: But, Matt, the reason we are so impressed is that if it were a classical particle there would be no chance of its getting through that barrier.

DR. MAYNARD: A classical point particle, yes, absolutely right.

MR. WAKELAND: A potential barrier higher than it is. See, this is like a ball coming to the wall and there would be no chance of getting through it. In quantum mechanics, there is some chance of the "it" getting through.

DR. MAYNARD: Classically, the velocity of the particle here would have to be imaginary, so it could never have a velocity in here, it could never get inside the barrier. That is a classical point particle.

That is only an approximation. Let's see what happens in the real world. [Computer demonstration]

Taking the modulus of the wavefunction squared -- remember, I said when you do a measurement you are really measuring an expectation value of some operator, but if the operator is the identity operator, then you are taking the modulus of the wavefunction. A lot of things in quantum mechanics involve the modulus of the wavefunction.

I really think that is not a great way of looking at things. You really should think of things in terms of measurements.

Here is the modulus of that wavefunction. As an acoustician, I tend to show the real part of wavefunction but, in quantum mechanics, you probably should show the modulus of the wavefunction. Here is the modulus of that wavefunction now and watch what happens when it hits the barrier. There is the envelope of that one I just showed. Watch what happens when this guy hits the barrier. I did not expect this; I was surprised. I mean, once you see it, well, of course.

Remember, you have an incident wave packet, but when it hits the barrier there is a reflected wave packet and they are interfering, so you get a standing wave. Also notice here, here is the part that tunneled through the barrier. There is the quantum mechanical particle tunneling through a barrier.

[Transparency 35]

This is the square well potential -- we just did the square barrier potential -- this is the square well. The point is here, if E is above here, then you just get waves everywhere but you can get interesting transmission resonances. That is very interesting. I should have a demonstration of that, but I do not.

There is the other case, where the energy is down here, and then you actually get what are called bound states, where you get discrete eigenvalues now, finally.

When it is like this over here, you have evanescent waves decaying that way, over here you have evanescent waves decaying that way, inside you have waves waving, and you have to match the boundary conditions and, in order to do that, you find that for this condition you can have solutions for only discrete values of E .

[Transparency 36]

Here are the solutions. The dashed line is for the finite value of V_0 . Here is an eigenfunction, this dashed line, here is one. You get symmetric, anti-symmetric, symmetric eigenfunctions. If you let V_0 go to infinity, then it is just like a string clamped at these positions, but we all know about that. The thing is, if you have a finite V_0 , you get these bound states; you get discrete eigenvalues at these values here.

[Transparency 37]

Let's do the simple harmonic oscillator. Simple harmonic oscillator potential is $1/2Kx^2$, or $1/2m\omega_0^2$ times x^2 , where K is the spring constant, m is the mass, and ω_0 is the classical oscillator's natural frequency.

But now we take this potential and put it into the Schrodinger equation and we get this. This is an ordinary differential equation that you have to solve for ψ and eigenvalue E . This is the outline of how you do it. You make some transformations; you get this equation. You can find solutions to this equation only if E is equal to an integer n plus a half times $\hbar\omega_0$. Otherwise, the solution blows up. If you have these discrete energy eigenvalues, or the discrete values for this parameter E , then the solutions are hermite polynomials; they do not blow up.

[Transparency 38]

Here is what it looks like. This is a normalized wavefunction. Remember that the inner product of the wavefunction with itself has to be 1, so you get some constants here.

Here are the eigenfunctions and the eigenvalues, energy eigenvalues, for the simple harmonic oscillator. For $n = 0$, you actually get a finite value for energy $\hbar\omega_0/2$. This

does not have a classical analog, but when you do a measurement, things work out fine. Here are the hermite polynomials and here are the sorts of wavefunctions that go with those.

Classically, you cannot have a negative kinetic energy, so in here E is bigger than V_0 and the difference is the kinetic energy, so you have classical turning points -- right here, right? A classical particle cannot beyond there. At higher energy, the classical turning points are here, for example.

Of course, classically, if you have more energy, the amplitude goes as the squared of energy, so the more energy you have, the more amplitude you have in the classical oscillator. Those are the classical turning points.

The wavefunctions have this exponentially decreasing part. It goes as the $-x^2/2$, so when it gets beyond the classical turning point you start seeing the exponential decay of the wavefunction. For the quantum and the classical, there is some correspondence there.

But now let's take a look at the actual wavefunctions, and so on. [Computer demonstration]

[Transparencies 39-44]

This is the simple harmonic oscillator. This is like the classical oscillator down here, mass on a spring, and this is the wavefunction for the case $n=10$.

Now let's turn on time and see what happens. The mass goes back and forth and the period is $2\pi/\omega_0$. If you look at the wavefunction, it is oscillating many times in the time that it takes the mass to go back and forth. There is something wrong here, there is no correspondence between the classical and the quantum.

The reason is, for the quantum mechanical wavefunction the time dependence is $E(iE/\hbar)$, which is $n+1/2$ times ω_0 times time. The wavefunction is oscillating 10-plus-a-half times faster than ω_0 . This wavefunction by itself is not a good representation of a mass on a spring.

Now instead of looking at the wavefunction, or its real part, let's take a look at the modulus squared. There is the modulus squared of the wavefunction and let's see how that corresponds to the classical mass on a spring.

That is not so good, either, because the mass on the spring is going back and forth but the modulus squared is just sitting there. As I said, if you solve just the boundary-barrier problem, nothing happens. That is what you get.

If you do a measurement -- I probably will not get to perturbation theory -- if you do a measurement, you find that the measurement couples 2 states together. Let's take a look at what happens if we couple 2 states. We will assume that the measurement couples state $n=10$ to state $n=11$. Let's turn on time.

If I did a measurement, I would see the mass on the spring going back and forth. In fact, if I made a wave packet you would really see the wave packet going back and forth just like the classical particle.

MR. POESE: I do not understand why there have to be 2 coupled states in order to see that.

DR. MAYNARD: I mean, that is just what the calculation shows. The modulus just sits there, the actual eigenfunction oscillates too fast.

DR. MIGLIORI: If you want to change the state with a harmonic oscillator, meaning you want to pump energy into it, you have to go from one allowed state to the other; therefore, you are external. The only thing you have access to externally is coupling those 2 states somehow. Therefore, the only frequency that counts is that beat.

DR. MAYNARD: I will show that in another viewgraph.

Some people also claim that you will not see classical physics until you go to high quantum numbers, so let's go all the way down to $n=0$ and see what happens if you make a measurement that couples $n=0$ to $n=1$. [Computer demonstration]

On the previous one, you will notice that the wavefunction went from one side way over to the other side and the classical particle at that energy had a large amplitude.

Here we are at $n=0$ and $n=1$ coupled together and you will notice that at that energy there is a smaller amplitude. If you make a measurement, everything works out just fine. It oscillates back and forth at ω_0 , even when you have small quantum numbers.

[Transparency 45]

Here is the viewgraph that does the math for what I just showed. Classically, you have positioned as $A \cos(\omega_0 t)$ the energy as $\frac{1}{2} m A^2 \omega_0^2$, so you can write it this way. If you take the quantum expectation value of position between the same states, you get 0. If you make a measurement that couples 2 states, let's take a look at the expectation value for the superposition, or just go through the math.

Again, one state has $n + \frac{1}{2}$, the other state has $n + \frac{3}{2}$, so the product, with the complex conjugate, is your $E(i \omega_0 t)$. You take the real part and you get $\cos \omega_0 t$, the classical frequency.

The amplitude, you can do this integral with the hermite polynomials and you find that is exactly the same amplitude that you get for the classical part, but you get that only if you couple 2 states together.

Also, this is nice for these quantized sound waves. What a phonon is, is not some state with quantum number n , a phonon is when the system goes through a transition from n to $n +$ or -1 .

MR. POESE: A phonon?

DR. MAYNARD: A phonon, quantized sound waves in a solid.

MR. POESE: It is only defined, or whatever, if you are changing states, if you are adding energy or taking --

DR. MAYNARD: Yes, if I had a solid, if I had just 1 eigenstate, it would just sit there and I would see nothing. If I do a measurement, it is going to couple 2 states and then I see the quantized sound wave.

In fact, you can make a wave packet. People in atomic physics are taking high-order n states and actually making electron wave packets that go around classically. That is another story.

DR. FORTUNKO: You mixed those 2 wavefunctions on a one-to-one basis, it was the sum of the 2 wavefunctions?

DR. MAYNARD: Yes, maybe there should have been a squared of 2 or something in there.

DR. FORTUNKO: No, no, I mean is there a reason for that? Can you mix them in other proportions?

DR. MAYNARD: I had not really thought about that. You have to really consider what a measurement is doing and, in fact, I have got viewgraphs for perturbation theory that -- I do not know, my hour-and-a-half is up.

DR. FORTUNKO: Well, keep going, keep going.

[Transparency 46]

DR. MAYNARD: Let me quickly flash what comes up. There are central forces that use spherical coordinates. For classical central forces angle of momentum is conserved.

[Transparency 47]

Here is the quantum version of central forces. You get the Laplacian in spherical coordinates. It is a differential equation that you can solve, separation of variables. What comes out, the theta and phi parts, the eigenfunctions are spherical harmonics.

[Transparency 48]

This is quantum angular momentum. If you crank it all out, you get this. If you look at L^2 , it is just the same thing you get for the Laplacian in the differential equation. The Hamiltonian you can write like this. Now all the angular stuff is in the L^2 . This looks a lot like the classical thing with a centrifugal force term.

From all this manipulation, you find that L^2 commutes with the Hamiltonian and the Z component of angular momentum commutes with the Hamiltonian. If you go back to the equations that means their expectation values are constants in time and, in fact, they are also eigenvalue operators for the spherical harmonics.

You could spend months in a quantum class in dealing with spherical harmonics and angular momentum, and the addition of angular momentum. It is a nightmare. Klepps-Gordon coefficients -- does anybody remember what Klepps-Gordon coefficients are?

Okay, we are not going to do that.

[Transparency 49]

Then you can actually put in a Coulomb central force and get the hydrogen atom details. You can read about it.

[Transparency 50]

Here is a nice picture. These are the radial functions. As you go to higher-order radial quantum numbers, you see the average radius goes out, which is just what you would expect classically. These are the polar plots of orbital-to-angular momentum theta and phi dependence.

[Transparency 51]

I am just going to mention these things now. Historically, people put atoms through an inhomogeneous magnetic field and they split. That means there is an angular momentum but the actual angular momentum for these atoms was zero.

Dirac did a theory that extended the classical Hamiltonian theory that we did to include relativistic dynamics, so you got the square root of $1 + V^2/C^2$ sorts of things. This theory predicts that the electron has an intrinsic angular momentum that is called spin.

To these wavefunctions that we have been talking about you actually have to add 2-element vectors. The spin has only 2 states, minus a half and plus a half, so you have 2-element vectors

that are called spinners. The operators are 2×2 matrices, which can be written as a linear combination of Pauli spin matrices. You can read about that.

[Transparency 52]

This is really interesting. Let me do a little more with this. For systems with many particles, the Hamiltonian will have generalized coordinates for each particle. For classical identical particles, the particles can be distinguished by following their precise classical trajectories.

In quantum mechanics, particles are waves. Because of the principle of superposition, when the quantum particles overlap, you cannot distinguish them or just add up, you cannot say here is one and here is the other, so quantum particles are indistinguishable and that has some pretty spectacular results.

If you have identical particles, then the wavefunction can be written as a product -- like this. If they are indistinguishable, if I switch the quantum numbers for 2 different particles, I should get the same state, because they are indistinguishable. This is called exchange degeneracy.

There are only 2 wavefunctions where ψ^2 is invariant under exchange, and that is one that is symmetric and one that is anti-symmetric. I am really going through this fast.

[Transparency 53]

We will stop here, I think. If you take 2 particles, 2 sets of quantum numbers, and switch them, in the symmetric wavefunction they do this, but in the anti-symmetric wavefunction, if you switch them, a minus sign pops in.

If you have 2 particles that have the same sets of quantum numbers, K_i and K_i , this is K_i and K_i , that is the same function now and the only way a function can equal to minus itself is if the function is 0. For the anti-symmetric wave function, you cannot have 2 particles having the same quantum number.

Historically, from experiments on multi-electron atoms, people noticed, or Pauli expressed (the Pauli exclusion principle), there can never be one electron in the same quantum state. What that means, if you take a look at this and that together, is that electrons must have anti-symmetric wavefunctions.

The generalization is that particles with half-integer spin quantum numbers must have anti-symmetric wavefunctions, and these particles are fermions. Particles with integer spin quantum numbers must have symmetric wavefunctions and they are called bosons.

If you do statistical physics with these 2 restrictions, that these fermions cannot have the same set of quantum numbers, when you do statistics you find that fermions have to obey Fermi statistics with a distribution function. In other words, the number of particles that have energy E is given by this expression. For bosons, the number of particles that can have energy E is given by this expression.

This one can never be bigger than 1. This one, because of the -1 here, can actually blow up. If you go to low temperatures, you find that bosons can all pile up into the same energy state, and that is Bose condensation.

Transparencies 54-55)

The other viewgraphs I had were perturbation theory and time-independent perturbation theory. When you make a measurement, you are basically doing time-dependent perturbation

therapy, so this is what you have to know when you are doing measurements. If you do everything carefully, everything works out fine. You can read about that.

The end. Thank you.

PERIODIC, RANDOM, AND QUASIPERIODIC MEDIA

Julian D. Maynard
Department of Physics
Pennsylvania State University

DR. MAYNARD: Thanks.

[Transparency 1]

Here is an integrated circuit. It is a CPU for a computer. You might ask what does this have to do with acoustics. Those of you who use Mathematica and do computer simulations may think this is acoustics. (Laughter)

The goal with integrated circuits is to make them go faster and faster and, in order to do that you need to make the components, the transistors, and the wires that connect the transistors smaller and smaller. So you need to ask: does the physics for the larger versions, the regular size version transistors, still work for the smaller versions? For example, for the small connecting wires, is Ohm's law still valid?

Suppose you had a wire that was only a micron in length and only 100 Å across, you put current through it, and you measure a voltage drop and check to see if the voltage is proportional to the current. For a while like this you would find that, in fact, Ohm's law is no longer valid.

Microscopically this wire has electrons and positive ions in it, so to understand that system you have to deal with a fundamental problem in condensed matter physics. That problem is to solve the Schrodinger wave equation, which I talked about on Sunday, for an electron moving in the potential field of some 10^{23} ions. Well, 10^{23} is a big number, so even with Mathematica you would probably have a hard time solving that problem.

The electron may be in a solid crystal and the ions are arranged periodically in a crystal. It turns out that you have to solve the Schrodinger wave equation for only one unit cell, and that is something you can do. Electrical engineers have computer programs that they can use to calculate band structure of semiconductors, and so on.

Unfortunately, crystals are not perfect; they have defects in them. In fact, there are some very important materials, like amorphous materials and alloys that can be highly disordered. With these highly disordered materials you are back to having to solve the Schrodinger wave equation for some 10^{23} nonperiodic scatterers.

Fortunately, also, in the real world you can make measurements at finite temperatures. At finite temperatures, these ions are moving and, in fact, these moving ions are lattice vibrations in a solid, those phonons, the quantized sound waves that I told you about on Sunday.

An electron can inelastically scatter off of one of these phonons and that inelastic scattering sort of destroys the wave nature of the electron, destroys its phase coherence. Because of that inelastic scattering and this randomization, we can assume that certain things average to 0 and you can use statistical physics.

With statistics you can convert the wave equation with some 10^{23} parameters into a transport equation, a Boltzmann equation or diffusion equation, with just a few parameters; in fact, maybe just one mean free path for this inelastic scattering. Because of this inelastic

scattering and the randomization, solid-state physics has been able to continue and deal with disordered materials.

Relatively recently a couple of things have happened that have changed all that. First of all, the experimentalists made their samples so small and so cold that the size of the sample was smaller than the mean free path for this inelastic scattering, so that essentially gets rid of the inelastic scattering.

Without the inelastic scattering, these statistical techniques are no longer valid and you are back to having to solve the Schrodinger wave equation for some 10^{23} nonperiodic scatterers, so it is a big problem.

[Transparency 2]

It is a very fundamental problem. It is essentially wave propagation in a disordered array of scatterers and that is a problem that dates back to Lord Rayleigh. Since Lord Rayleigh, a lot of famous people have worked on this problem. But because of these more recent experiments with these small cold samples more theoreticians have gotten involved and they have actually been able to make some progress for wave propagation in disordered arrays of scatterers.

These contemporary theoreticians did not want to say that they were solving the wave equation, because that is what Lord Rayleigh did. Instead, they say that they are studying long-range phase coherence. And in fact, they are not actually solving the wave equation with a large number of disordered scatterers.

What they do is they go back to the equation that they can solve, the diffusion equation, and then they do perturbation theory back toward the original wave equation, so they try to take the diffusion equation and then patch it up with as much phase coherence as they can. They have quasi-particles called diffusons, and so on. Anyhow, they were able to make some progress.

This became a big frontier area in condensed-matter physics. There were papers published on Anderson localization, which I will be telling you about, universal conductance fluctuations, superdiffusion, nominal electron persistent currents, thousands of papers published.

I said that 2 things happened. The second thing that happened was that quasi-crystals were discovered. These quasi-crystals -- it is an aluminum alloy, it is made by rapid quenching, and the original samples were quite small, only a few 100 μ in size. I will tell you more about quasi-crystals later in the talk.

For now, I want to go back to these small cold samples. The actual experiments that people did were they made small wires like this, 1 μ long and only 100 \AA across, and they measured electroconductivity in these small wires. In physics there was a name given to this area of research. It was called mesoscopic physics.

[Transparency 3]

It means the study of phase coherence on the scale of microns. For our research we coined our own phrase and that is "megascopic physics," so megascopic physics is studying phase coherence of a scale of millions of microns. But of course a million microns is a meter, so we like to claim that we can see phase coherence in a 1-dimensional wire that is not a micron long but 10 m long.

We study properties of quasi-crystals that are not only a few 100 μ in size but more than a meter in diameter, certainly the world's largest quasi-crystal.

These statements are worded so as to sort of shock people whom actually work in these areas. Our 1-dimensional wire is not made with electron beams; our 1-dimensional wire is actually a piece of music wire that we bought at the music store that you use on a guitar. Our quasi-crystal is made out of aluminum alloy but it is the kind of aluminum alloy you find in the machine shop, not made by some special quenching method.

Our experiments are classical analog systems. In particular, they are acoustic analogs. One way that you can think about them is that they are analog computers that are used to solve problems in quantum mechanics or mesoscopic physics.

These acoustic analogs have some very nice advantages. First of all, they are very precise analogs of the quantum mechanical system. Let me go into that in a little more detail.

[Transparency 4 – unavailable at time of printing]

To learn about the properties of solids, you need to use quantum mechanics and I talked about this on Sunday. You need to solve the Schrodinger wave equation, which is shown here, and you need to find the solution ψ that describes the behavior of particle energy E , which is acted upon by forces that have a potential field, which I wrote as V/R here.

The solutions are eigenfunctions that may exist for only discrete values of the energy E for the eigenvalue, so you get quantized energy levels. The thing is, you can take this equation and just multiply through by $-2/H \text{ bar}^2$ and you will get this equation, where the Q here is the square root of $2ME/H \text{ bar}^2$.

Now take a look at the acoustic wave equation. It looks like this. Here, in the Schrodinger equation, if you had the time-dependent Schrodinger wave equation, the time dependence here would be $e(i)$, or $-i$, ωT , and ω was $E/H \text{ bar}$. For the acoustic system you can assume a time dependence $e(-i \omega T)$ the same way.

Then you can also rig up an acoustic system where the speed of sound varies with position -- like this. If you plug this in, you just do a Taylor series expansion and you can rewrite the acoustic wave equation this way, where now this Q (sort of an eigenvalue parameter) is just ω over some nominal speed of sound.

In quantum mechanical systems, this Q is related to energy eigenvalues and in the acoustic system this Q is related to eigenfrequencies, or natural frequencies, for the system. If you look at these 2 things, these equations are mathematically identical, or at least you try to rig up an acoustic system that has the same sort of potential field as the quantum mechanical system.

An important thing is just the symmetry of a problem. The title of my talk is "Periodic, Disordered, and Quasi-Crystalline," (sic) and those just refer to the types of symmetry that the potential field can have. You can have acoustic systems that are mathematically identical to quantum mechanical systems; you can use them like analog computers.

[Transparency 5 – unavailable at time of printing]

Some other advantages to the acoustics systems: In an experiment, all the conditions can be very precisely controlled, or at least measured. They are megascopic, meters in size. When you do a measurement, you can measure eigenvalues, eigenfunctions, and other properties. For a system governed by a differential equation like this, once you have measured eigenvalues and eigenfunctions that is all you need to measure. Everything else can be calculated in terms of those quantities.

These first 3 things you can also do with digital computer simulations, but in these acoustic systems you can do more. You can make this potential field time-dependent and you can also turn up the amplitude and make them go nonlinear. These are 2 things that are very difficult to do with digital computer simulations and quite a bit easier to do with the acoustic analog systems.

[Transparency 6]

Before I show you the experiments, I want to explain how we got into this business. As you know, at Penn State, there is a lot of acoustics research in different areas, and one of the areas is noise reduction. Here is a typical problem in noise reduction.

Suppose you have a plate, like a floorboard in a small aircraft that flies into Monterey, and some machinery in there that creates vibrations in the plate, and those vibrations travel down the plate as a wave. As the plate vibrates at the other end, it radiates sound and makes some annoyance. Usually, for structural reasons, the plate will have a rib on it and a rib will reflect the vibrations so that less vibration gets transmitted, so you get less noise at the other end of the plate.

[Transparency 7]

The engineers at Boeing Aircraft said, well, if one rib reflects the vibrations, why not a whole bunch of ribs? So, either for ease of manufacture or aesthetic reasons, they put down a nice periodic array of identical ribs. What they found is that the first couple of ribs reflected the vibrations but all the rest transmitted without any further reduction at certain frequencies.

This is something you learn about in solid-state physics or quantum mechanics. It is analogous to an electron in a metallic crystal. In a metallic crystal you have electrons and positive ions. If you measure electrical conductivity, you are trying to get the electron to move from one end to the other, but the electron is very strongly scattered by a positive ion. As it tries to go down, it gets turned around by these positive ions. It is very strong scattering.

From the strong scattering you would have to conclude that a metal should have very low electrical conductivity, the electron keeps getting deflected from its forward motion by the positive ions. Of course, metals have very high electrical conductivity, and there are 2 reasons. One reason is the electron is not a particle, it is a wave, it obeys the Schrodinger wave equation. The other reason is that in a crystal the positive ions are arranged periodically. The fact that it can move freely and that the waves are transmitted here are due to the fact that the scatterers are arranged periodically.

[Transparency 8]

Both of these effects are a consequence of a theorem that you can prove using group theory. In mathematics, it is known as Floquet's theorem but, of course, in physics we cannot have our theorems named after mathematicians, so we had to rename it Bloch's theorem.

What Bloch's theorem says is that if you have a system with a periodic potential or, in acoustics, if you have a periodic impedance, the eigenfunctions are said to be extended; that is, they have this particular form here. There is e^{ikx} phase factor that looks like a plane wave, but then it is multiplying a function and that function is periodic with the period of the potential field.

Now, if you take the modulus of this wave function, then this phase factor goes away and all you are left with is this periodic function. If you evaluate the modulus at the same point in

any unit cell, you will get the same value. In other words, what I say is that the wave function has the same nominal amplitude throughout the entire system.

Over here, I have copied some Bloch wave functions from a textbook. This is Michael Tinkum's book on group theory, a very good book. Here is a periodic function, but with the phase factor. Even though it is modulated, if you look anywhere in here it will have the same nominal amplitude.

It is the same amplitude throughout an infinite system.

In quantum mechanics, the wave function modulus gives you the probability of finding the electron at that point. If it is an extended wave function, then the electron has equal probability of being in any unit cell or, in other words, it can travel freely through the whole system. That is why metals have a very high electrical conductivity.

In solid-state physics, Bloch's theorem is kind of taken for granted. It is actually a very subtle effect. If you look at it carefully, it is also quite complicated. If you ever deal with a periodic system, say a periodic waveguide, take care to treat the Bloch eigenfunctions carefully.

[Transparency 9]

This viewgraph restates the situation. If you have a potential field that is periodic, then the eigenfunctions are extended, the same nominal amplitude throughout the system. Now you can ask what happens if you take the periodic potential field and add a part to it that is random.

What you might think -- what I thought -- is that if you add just a random part, then you would just get some random fluctuations in the amplitude of the wave function, maybe go up here, down here, up, and so on. It turns out that that is not what happens. Yes, question?

MR. SMITH: On the upper graph, if it is supposed to be a periodic variation, what happened on that far right, that last peak? Shouldn't that be dropped?

DR. MAYNARD: It depends on the value of this k here. This, in fact, is called the Bloch wave number. As I say, it is a subtle effect. It might skip 2 unit cells.

What happens if you add disorder? What happens is the eigenfunctions become exponentially localized; that is, there will be some point in the whole system where it has a maximum value and then, as you go away from that site, the eigenfunction exponentially decays. That is something I would not have expected. It is a very nontrivial result of the disorder.

In fact, Phil Anderson and Sir Neville Mott used this effect to explain the metal-to-insulator transition in disordered metals. The level of sophistication here in this effect is worthy of a Nobel Prize. It is not a trivial matter. This effect is called Anderson localization and the characteristic length -- remember, this is exponentially decaying, so there is some characteristic length of that exponential decay and that is called the Anderson localization length.

As I mentioned already, this is a very fundamental problem, it is wave propagation in a disordered array of scatterers. Lord Rayleigh worked on that problem. Someone told me that Lord Rayleigh almost got Anderson localization, but he did not have Mathematica, so he had to get his Nobel Prize for a different reason.

Since Lord Rayleigh, a lot of famous people worked on the problem, Wigner, von Neumann, Freeman Dyson. They actually made a lot of progress in understanding waves in disordered media.

The first rigorous theorem on waves in disordered media is Furstenberg's theorem and that was proved in 1963. That is quite recent on the time scale of Lord Rayleigh. This is a

fundamental problem, it has been around for a long time, and it is only relatively recently that anyone has been able to prove a rigorous theorem.

Not only that, Furstenberg's theorem is for only one dimension. In 2 and 3 dimensions there are no rigorous theorems yet. There are so-called scaling arguments but they are so-called arguments, they are not rigorous theorems.

When I first heard about Anderson localization, I was pretty amazed, because I would have bet money that the disorder would have just put disorder into the otherwise extended wave function, so I wanted to learn how disorder caused this exponential localization of the eigenfunctions.

I went to the library -- I should mention here that I am an experimentalist and there is a saying for experimentalist that 6 months in the lab can save you a day in the library -- I violated that saying and went to the library and spent quite a few days. There were thousands of papers written about Anderson localization, but all of those papers were just using Anderson localization as a way of explaining some observation in solid-state physics. They did not give you a really nice balls-and-springs explanation of why disorder gives you this exponential localization.

I tried to read Furstenberg's paper. Furstenberg's paper is 25 pages of theorems and lemmas and it references other papers with theorems and lemmas. Again, I am an experimentalist. I have a very hard time reading papers that start off with "let M be a manifold." (Laughter) I could not get past the first paragraph.

Actually, since that literature search I did find a paper by Marshall Luban and Jim Luscombe that gives a beautiful description of how disorder gives you Anderson localization.

Let me give you some idea of what the problem is in understanding this. The statement of Furstenberg's theorem is something you cannot understand. Furstenberg's theorem says that in a 1-dimensional infinite system if you have disorder, then the eigenfunctions will decay exponentially to 0 as you go to plus-and-minus infinity with probability 1.

I can understand those words, but as an experimentalist, I have some trouble with that, because, first of all, I do not have any infinite systems. I have to try to understand Anderson localization for a finite system. Even if I had an infinite system, I am supposed to go marching off to infinity and look for something happening with probability 1. Well, how would I know if I had seen it happening or not?

Marshall Luban's paper explains Anderson localization for a finite system and it has this probabilistic nature built into it, so if you want to understand Anderson localization, take note of that paper. I had to read the paper twice, but the light finally dawned.

DR. ATCHLEY: There is a loop here to UCLA. Jim Luscombe was a graduate student from UCLA and he worked with Keolian and Baker.

DR. MAYNARD: Yes, I know. I taught him in the 1NED lab.

DR. ATCHLEY: Now he is at the Postgraduate School.

DR. MAYNARD: Oh, is he? Super. He is a really nice guy.

DR. ATCHLEY: We will get him to do acoustics yet.

DR. MAYNARD: The Anderson localization is a consequence of the statistics of this disorder. In statistical physics, one of the most important parameters is the dimensionality of the

system. It turns out that in one dimension any amount of disorder will give you some degree of Anderson localization -- that is Furstenberg's theorem.

If the amount of disorder is small, then the Anderson localization length may be quite big. If you have finite system and the Anderson localization length is large, you may not see the effects of the localization if you have weak disorder. That is why systems, even though they are not perfect, can still look periodic. As you increase the amount of disorder, the Anderson localization length gets smaller and you can see the exponential localization of the eigenfunctions.

In one dimension, any amount of disorder will give you some degree of localization. In 3 dimensions there is a critical disorder that you have to exceed before you get to localized eigenstates. If the system has disorder that is below the critical disorder, then you can still have extended eigenfunctions, but if you exceed the amount of critical disorder, then you get the Anderson localized eigenstates. In fact, that is the explanation of the metal-to-insulator transition for disordered metals for which Anderson and Mott got the Nobel Prize.

Two dimensions is the so-called critical dimension. It is hard to do the theory in the critical dimension, so people were not sure if in 2 dimensions there was a critical disorder above which you get the localized states and below which you get extended states.

More recently, with the scaling theory, they think they understand it, but at least at the time it was a good puzzle. So we thought we would do an experiment to study 2-dimensional Anderson localization.

[Transparency 10]

Since we knew about the vibration of plates with ribs, the experiment we thought we would do is we would take an aluminum plate, maybe a couple of millimeters thick, but maybe meters on a side, and we could drive the plate with a shaker and excite transverse waves in the plate. That would be our 2-dimensional wave medium.

For a potential field, for an array of scatterers, we were going to put a rib structure on the plate. On one plate we would put a periodic rib structure, a periodic hexagonal rib structure, and on this plate, if we drive it and look at its normal modes, we should see the extended Bloch wave functions, because it is periodic.

Then we could construct another plate that had a disordered rib structure on it. The way you construct this is you take the centers of these hexagons and displace them by some random amount, like this, and then you do a Wigner site cell construction; that is, you connect the centers, the new centers, and then you put on perpendicular bisectors and those become the boundaries of these polygons. So, this is a disordered array of polygons.

MR. SMITH: This is a purely theoretical construct you are making?

DR. MAYNARD: No, we were actually going to make plates that had these rib structures.

MR. SMITH: I would hate to be the machinist on that.

DR. MAYNARD: Well, they have computer-controlled machines.

If we drive this plate and look at its normal modes of vibration, we should see Anderson localized states.

Before we did this experiment, we thought we ought to do a 1-dimensional experiment first just to educate ourselves as to how this Anderson localization worked. This is before I found the paper by Marshall Luban and Jim Luscombe, so it was still a mystery as to how disorder

produces this exponential localization and I was not sure I believed it, so I wanted to see it happening.

[Transparency 11]

For our 1-dimensional experiment, we went down to the local music store and bought a spool of music wire -- wire like you use on a guitar. In fact, we found that harpsichord wire has the highest Q . We took about 10 m of this steel wire and hung it down the stairwell in the physics building, with a weight at the bottom so that there was a tension in the wire.

Then we took a shaker and put it next to the wire so this would excite transverse waves in the wire. This is just like the vibration of a guitar string. That is our 1-dimensional wave medium. For the potential field we took little lead masses -- in fact, these were split shot like you use for fishing lines -- and put them along the wire. These are the scatterers. These would be just like the positive ions in a metal.

The masses on a wire act like a delta function. This system is a very accurate analog of a 1-dimensional wave equation with a Kronig-Penny potential field (that is something you study in solid-state or quantum mechanics classes). We could put the masses on here periodically and look for the extended Bloch wave functions, but then we could put them on disordered and look for the Anderson localization.

Running parallel to the wire there was an aluminum beam with a trolley on it and a magnet. This magnet put a magnetic field locally across the wire. When the wire vibrated in the magnetic field there was a current induced in the wire, so we grounded one end and put the other end to a current amplifier.

From the amplitude and the phase of the current here you could determine the amplitude and phase of the vibration of the wire at this point. We could translate the magnet down the wire and we could make a plot of the amplitude and phase of the vibration of the wire. In other words, we could measure the eigenfunction, amplitude, and phase directly. That is something that is virtually impossible to do in quantum mechanics.

With these scanning small probe systems people are trying to measure quantum mechanical wave functions but no one has succeeded yet.

Now we could measure eigenfunctions directly with this system for either a periodic or a disordered system, but before you measure eigenfunctions you need to know the eigenfrequencies or the eigenvalues. To find the eigenvalues we would just put the magnet in one position and then drive the system and sweep the frequency. When the frequency corresponded to one of the eigenfrequencies, we got a resonance in the response.

[Transparency 12]

Before I show you those data, I want to show what you would expect to see from theory. There are 2 ways of trying to understand this problem. One is the free-electron approach and the other is the tight binding approach. The easiest way is this tight binding approach.

Here I have drawn the wire horizontally. You first imagine that the masses are infinite. If the masses are infinite, it is as though the wire were clamped at these positions. If the wire were clamped here, then it is easy to understand the eigenfunctions, the normal modes and the natural frequencies. They just correspond to fitting an integral number of wavelengths between the clamps, like one half-wavelength or 2 half-wavelengths.

The frequencies, there is some fundamental and the other frequencies are just integer multiples of that fundamental frequency, so you get a harmonic series. This is if the masses were infinite and, of course, the masses are not infinite, they are finite, so when one section of the wire here vibrates, then the mass on either side vibrates, so it can be coupled to the neighboring sections of wire.

One nice way of thinking about this system is you can think of the section of wire as a local oscillator and by itself, as an isolated local oscillator, it has these sharp eigenfrequencies. With the finite masses, each local oscillator is coupled to its neighbors and when you take a system of local oscillators and couple them together the sharp frequencies of the local oscillators broaden out into bands. This is a very nice demonstration of how you get band structure in solid-state physics.

Between the bands, there are gaps where you actually have evanescent waves in the system, but essentially, there is very little response of the system in these gaps. In solid-state physics these would be bands and gaps and in an acoustic or microwave system these would be pass bands and stop bands.

I have a short version of this with just a couple of masses on it and you can demonstrate band structure, just a benchtop experiment.

The system is finite in size here. We had about 50 masses on the string. With 50 discrete masses, that means that each one of these bands really has only 50 discrete eigenstates in it, so each band has 50 eigenfrequencies like this.

MR. SMITH: I am trying to get a sense of scale from one to the other. For example, where would ω_0 appear on the scale?

DR. MAYNARD: There is a mode here. There is another mode down here that corresponds to just the whole thing being translated. When you add the coupling, which goes into modes that have long wavelengths like this, all the masses are moving to make one big wavelength.

This one corresponds to approximately to one half-wavelength between the masses, so you can have a state that looks like this, where this bit right here pulls the mass a little bit, so you get the coupling there. You can really go into Bloch's theorem in great detail. All this is for the periodic system that I am showing you now. It is actually quite complicated.

This one up here is this second band here. All of these normal modes here would correspond to fitting approximately one half-wavelength and this would be fitting approximately 2 half-wavelengths.

Now you can see what happens if you sweep the frequency. At a certain frequency you can start out in a gap here. Then as you sweep the frequency you will go into a band and when you go into a band, then you get these 50 resonances, then into another gap, where you get no response, then to another band, and so on.

[Transparency 13]

Now I can show you the actual data. This is the response of the system as a function of frequency. Here you get virtually no response, so this is in a gap. Then, as you go into a band you get about 50 resonances. As you go higher, you go into another gap. As you go up quite a bit higher, you will go into the next band.

PARTICIPANT: What is the additional structure?

DR. MAYNARD: If you counted these, there would be about 50 there. The string with 50 masses on it has 50 normal modes in each band. This extra structure here is kind of interesting. The textbook treatment of periodic systems uses periodic boundary conditions and there are the eigenfunctions.

For periodic boundary conditions you get something that looks like this, and these are actually Bloch traveling waves. In fact, this thing here looks just like a plane wave, traveling wave.

You can have linear combinations of these that are Bloch standing waves. For a real system with finite boundary conditions, you get Bloch standing waves. It turned out that at these frequencies the pickup -- that magnet just happened to be close to a node in the Bloch standing wave, so the response was small here. We should have just moved it around to get a higher response at all the frequencies. That explains this structure here.

These are the eigenfrequencies for the periodic system. Now you can sit on one of these frequencies and then translate the magnet and plot the eigenfunction that goes with that eigenfrequency.

[Transparency 14]

This is a plot of the amplitude of the response as a function of position along the wire, so it is a plot of the amplitude of the eigenfunction. These are 2 different eigenfunctions for 2 different eigenfrequencies in the system.

They have some interesting structure in here. Again, this is something that you do not see in the solid-state textbooks -- you can get sort of nontrivial structure here. This is the sort of thing you might expect to see for a Bloch standing wave.

The masses are actually quite strong perturbations and the masses do not like to move very much. These positions down here are the positions of the masses, so there are about 50 of those approximate nodes in the wave function. You can see that there is about one half-wavelength between the masses. These are eigenfunctions in that second band.

These are pretty nice extended Bloch wave functions. It is pretty amazing, because one of the first things we did was to put one mass on the wire and sent in a wave and measured the reflection coefficient. The reflection coefficient was about 99.97%, so these are very strong scatterers. Yet, if you put them down periodically, then the wave goes right through the system, so these are Bloch wave functions.

One thing you might notice, also, is that there is a little structure here, they are not perfectly periodic. That is because the masses in the system -- these were just commercial split shot and the mass actually varied by about 13%, but the mass does not make so much difference. What is critical here is the frequency of the local oscillator, which depends on the spacing between the masses. That had to be set very, very carefully -- in fact, we had a pair of calipers that we used to set the spacing between the masses so that each local oscillator had the same frequency and then you get the extended Bloch wave functions.

PARTICIPANT: What about the tension of the wire?

DR. MAYNARD: While it was hanging?

PARTICIPANT: Yes.

DR. MAYNARD: We worried about that, too, but from the measurement we just determined that it did not seem to matter. You can probably go back, do the theory, and find that it does not matter, but that was an interesting problem to think about.

This is for the periodic system and it was important to verify that we could see the Bloch wave functions, because that is a rigorous theory.

Then we went back and changed the spacing between the masses -- they were about this far apart -- by a random amount but within 2%. We did a random number generator and then calculated a shift that had a maximum value of 2% of the lattice constant. When you do the disordered system with 2% disorder, there is a dramatic change in the eigenvalue spectrum.

What you find is that the eigenfrequencies tend to bunch up. There is one here, there is an eigenfrequency that appeared here, which would have been in the gap for the periodic system. This is the sort of thing that you hear people in solid-state physics talking about, states appearing in the gap. Here you can actually see it happening.

There is also a theorem that says if you have a state that is off by itself, kind of isolated, that state will be highly localized. If you have a state that has neighboring states around it, then that state will be less localized. We can actually see that in the experiment.

For the disordered system, we can pick some eigenfrequency and then translate the magnet and look at the eigenfunction. These plots here, except for this last one, these 5 plots are different eigenfunctions for different eigenvalues for the 2% disordered system and here you can see we got textbook Anderson localization for the state. So, it really did work. This is the first direct observation of an Anderson localized eigenstate. All the experiments on these mesoscopic electronic systems, you could only infer that the electrons were localized, you did not really measure it directly.

One thing you can see in these 5 different eigenstates here is that the Anderson localization length is different for different states. In fact, the transmission from one end to the other is quite different for different states and that corresponds to universal conductance fluctuations that people are studying with the mesoscopic systems.

MR. HICKEY: Is that the whole length of the wire?

DR. MAYNARD: Pretty much. There were 50 masses, but there was some more wire beyond that.

MR. HICKEY: So this band of localization has several masses?

DR. MAYNARD: Yes, the localization length would be about 6 lattice constants in that case.

MR. POESE: It seems like you could make one of these with a standing wave tube, also. You could put little masses in the tube to be constrictions to increase the local acoustic mass.

DR. MAYNARD: Yes, in fact, my first idea was to have a tube with little Helmholtz resonator side branches.

MR. POESE: And if you did that in an acoustic engine, it seems as though you could make a picture that looks like C so that you have all of the acoustic fluctuation and the stack, for example, and none anywhere else in the tube to reduce your viscous losses along the tube. Is that a possibility or is that silly?

DR. MAYNARD: I do not know. I will have to think about it. Yes?

MR. WAKELAND: When you say the masses do not move, there is this other experiment where you assume that there is no mass in the string, and --

DR. MAYNARD: Yes, the free electron approach.

MR. WAKELAND: But I am saying you would have the whole string moving off to the side one way or the other. We are above all of those frequencies here.

DR. MAYNARD: Yes, we are above that frequency. Mostly we are working in that band that corresponds to fitting approximately one half-wavelength.

MR. WAKELAND: So you are talking about totally different frequencies, all right.

DR. MAYNARD: Yes that is true, this would be a much higher frequency.

MR. WAKELAND: You are talking now about having several wavelengths across the stack.

DR. MAYNARD: Okay, so we could see the Anderson localized states, we could see the universal conductance fluctuations, and so on.

We did this experiment just to educate ourselves as to how this Anderson localization worked and that turned out fine. We had no intention of publishing this, but when we did the experiment, we realized that we could do some serious physics with this experiment and it went like this.

Suppose that at an eigenfrequency of 800 Hz you get the energy localized here. Then suppose that at an eigenfrequency of 810 Hz you get the amplitude localized at a different site.

Oh, that is another thing. Different eigenstates and different eigenfrequencies, the eigenstates are localized at different positions, different sites.

MR. SMITH: Does that vary a lot depending on your --

DR. MAYNARD: It depends on the actual realization of the disorder, yes, where they get localized.

We said, well, suppose we drive it at 800 Hz and get the energy localized here and then we modulate the tension in the wire at 10 Hz, the difference frequency. When you modulate the tension, the little masses are going to go like this a little bit and that looks just like a longitudinal phonon going through our system.

Remember, these transverse waves are analogous to the electron Schrodinger waves. If you modulate the tension so that the little scatterers are moving, that is like a longitudinal phonon. In this experiment, we could simulate electron phonon scattering -- in the second viewgraph, I mentioned the inelastic scattering between electrons and phonons. Here we could actually simulate it in gory detail in this megascopic acoustic system.

This inelastic scattering between electrons and phonons for a disordered system was some serious physics that theoreticians were trying to do at the time and we could do this experiment and see it in gory detail.

When you do the experiment, what you see is the energy starts out localized at one site and then you turn on this modulation, send in the phonon, and what you see is the energy going back and forth between the 2 sites at the difference frequency. [Demonstration]

This is a demonstration you have probably seen. Here are 2 states. There are 2 pendula. You can think of them as independent pendula, you can think of their separate eigenstates or normal modes -- or normal mode -- just oscillates like this.

Then you couple them together; there is a knot here that couples the 2 pendula together. When they are coupled, there are still 2 normal modes. One normal mode is like this and, in fact, the coupling is not even involved, so this frequency is not shifted from the simple pendulum frequencies. This is like a symmetric wave function and this would be like an anti-symmetric mode. This is the other normal mode and in this case, the coupling does matter.

So those are 2 coupled modes and now suppose you start off with all the amplitude over on this side, but they are coupled, so this one starts moving. The interesting thing is that at some point this one gives up all the energy, so it should come to rest -- right there -- and now this one is swinging.

This one, then, will start putting the energy back, after a while this one will stop, and now all the energy is over here. The energy goes back and forth at the difference frequency of those 2 normal modes. You can just think of them as beating together and this is the beat frequency.

MR. SMITH: So you are treating one as starting at zero frequency or did you vary the length slightly to get your beat frequency?

DR. MAYNARD: It has to do with this knot. Like this, the knot does not -- you have the whole length here. The knot changes the frequency, so the frequency of this one is different from the other one. If you start it here, it goes back and forth between those 2 frequencies.

[Transparency 15 -- unavailable at time of printing]

In the disordered solid-state system, the electron starts out localized at one site. A phonon comes along and couples it to another state, so it does this, but then at some point in time the phonon will go running off to the thermal bath; in fact, this is a good demonstration of what thermal baths are all about. They randomly send in phonons and they randomly take out phonons in an electronic system.

Depending on when that phonon leaves, there is some probability that the energy gets left in the second site. This is phonon-assisted hopping between 2 Anderson localized sites. This was a hot topic in condensed-matter physics, so we wrote this up and published this paper. It was published in *Physical Review Letters*. It is interesting, because the whole experiment cost about \$40.00.

In the experiment, you could measure the hopping probability as a function of how hard you modulate the tension in the wire. Down here, where it is linear, you can use first-order time-dependent perturbation theory, but then as you go to larger amplitudes, you start seeing higher order effects -- in fact, the thing saturates. Once you see this, you can figure out why it does that in a nonlinear model.

MR. WAKELAND: Do you turn on the modulation for a little while and then turn it off?

DR. MAYNARD: I am sorry. In our experiment, if you leave the modulation on and you measure the relative amplitudes of the 2 eigenstates, then the relative amplitude gives you the hopping probability. That is just interpretation of quantum mechanics. You do not really care when the phonon arrives or leaves; it is just what is the relative amplitude when it is there. That gives you the probability that it will get left in one state or the other when it does leave.

[Transparency 16 -- unavailable at time of printing]

While we are on this topic, let me describe an experiment that we are doing now along the same lines. This has to do with what happens if that potential field that was disordered is now time-dependent, not like this nice modulation but time-dependent in a random way.

If you had some static realization of the disorder, then you will get Anderson localization at some site like that, but now suppose this potential field is time-dependent, and you can imagine a really pathological time dependence that would just take this disorder here and shift it over. If you did that slowly, then the localized eigenfunction would shift over with it.

The whole idea of this Anderson localization is that the electrons get stuck and you do not get conduction, so this would be an insulator, but now with some time dependence the electrons can start moving again. If you have some time dependence that is not quite random but has some correlation in it, then the electrons can move and they will move faster than they would move if they were just diffusing through a disordered system.

They are not going as fast as just propagating as ordinary waves but they are not stuck, which would happen in a static disordered system; they are going faster than diffusion and that is called superdiffusion.

MR. SMITH: What would happen if you used several solenoids and varied those? Would that be the same effect?

You could achieve a random length change point to point from the mass.

DR. MAYNARD: I will show you what we do and what the interesting problem is.

You want to be able to change the local speed of sound. Our varying potential field is because we vary the local speed of sound and you want to be able to do that in real time, while the wave is propagating.

DR. MOLDOVER: Before you get to the dynamics, back to the airplane wing. If those ribs are placed randomly, do they function to block transmission of noise?

DR. MAYNARD: That was the sales pitch for doing this. We were thinking that with the Anderson localization different frequencies get localized at different points on the plate and then you can use some sort of narrow-band active cancellation techniques, or something, for doing that, but you use the plate itself to filter the noise.

That may still be a viable thing; we just got off onto other things.

DR. MOLDOVER: That is unresolved, you are saying?

DR. MAYNARD: Yes. I am pretty sure you can do it, but it is just how practical is it, is it necessary. I do not know.

Now we want a system where the speed of sound changes with position and you want to vary that in time. One way you can have the speed of sound vary with position is to use, say, a trough of water and look at surface waves. For surface waves the speed of the wave is the square root of G , gravity, times the depth.

If you have the depth vary with position, then the speed of the wave varies with position, so that is one way you could do Anderson localization. Actually, people tried this and the problem is there is too much damping in the water, too much attenuation.

Using the harpsichord wire with the masses was a really high- Q system. You do not want much damping, because you want this long-range phase coherence in order to see these Anderson localization effects.

You could imagine having a disordered potential field by having depth vary like this and then you could imagine having a rubber membrane on the bottom with little motors that drive the membrane up and down. You can change the potential field as a function of time.

The problem with that is when you drive the bottom up and down you generate a wave on the surface and that is not allowed. In the theory -- some theoreticians had a problem with this -- you had to change the potential field but you have to maintain the normalization of the wave function (remember, normalization is modulus squared integrated over the system has to equal 1). You can change the potential but you have to maintain that normalization. Some theoreticians fail to do that.

In our experiment, if you want to see the effect of a change in potential on the wave, changing that potential field cannot just generate a wave, because that is going to spoil everything.

[Transparency 17 – unavailable at time of printing]

Here is how we change the local speed of sound without generating a wave. What we use is a metal strip, it is a beryllium-copper strip, like this -- this actually has a mass on the end -- it has some stiffness, so it will oscillate by itself like that without having to put tension in it. The strip is going to be our 1-dimensional wave medium.

Then along the strip, we put a pair of piezoelectric transducers and they are pulled so that if we put a voltage across the piezoelectric one of them extends and the other one contracts, so it bends the strip across its width this way. If you bend a strip of metal like that, it becomes a lot stiffer in that direction. If you look at a venetian blind, that is why the venetian blind does not sag, because they bend it across the other way.

So, by doing this, if you bend it this way, you change the speed of sound in this direction, but if you bend it this way, you do not generate a bending wave in the longitudinal direction. We checked this out; it works very well. We could take the strip, drive it with a shaker, and excite a resonance for the transverse waves, but we can take one of these piezoelectrics and excite it at the proper frequency but it does not excite that resonance.

So here is a way you can change the local speed of sound without generating a wave in the system. We have an 8-foot piece of this beryllium-copper strip with 40 of these little piezoelectric pairs, each one goes to a circuit card, and then it all goes to a computer.

The speed of the wave in this beryllium-copper strip is pretty slow -- the highest frequency would be about 1 kHz, but with the computer and all these cards you can easily change the local speed of sound at 100 kHz or something.

[Transparency 18 – unavailable at time of printing]

Here is a picture of it. It is surrounded by a plastic pipe to keep airflow from making the thing vibrate too much, there are 40 of those chips, and here are the cards. It is a fairly sophisticated setup. This whole thing was built by an undergraduate student (he is now at Caltech). I will stop at this point and we will take a break.

DR. MAYNARD: The original idea was to study 2-dimensional Anderson localization --

[Transparency 19]

-- with a plate like this, but we got derailed from this experiment because quasi-crystals were discovered. Not too long ago it was thought that solids had 2 basic forms, crystalline and amorphous.

In an amorphous solid, the atoms are randomly placed and the properties are homogeneous and isotropic, but a crystal is quite different. To make a crystal you take a unit cell and you repeat it periodically to fill all space. In order to fill all space without leaving any gaps

anywhere, you can use only certain shapes, there are only certain shapes that will fill all space. There are 14 shapes and these form the basis for the Bravais lattices and classical crystallography.

Because there are only certain shapes that can fill all space, there are certain rotational symmetries. In particular, five-fold rotational symmetry is not allowed. Here is a quotation from Kittel's *Introduction to Solid-State Physics*: "A five-fold axis of symmetry cannot exist in a lattice because it is not possible to fill all space with a connected array of pentagons." (You can see here that there are going to be gaps.)

In 1982, Shechtman, in what was then the National Bureau of Standards, now NIST, made an aluminum alloy by rapid quenching and he put it in his x-ray diffraction machine and found 10 spots in a circle. That means that you have a mirror plane and *five-fold* rotational symmetry.

He showed this to people and, of course, they just laughed at him. They said, "Ha, you can't have *five-fold* rotational symmetry." So, he made better samples, other people made better samples and made measurements and, with better samples, these spots just kept getting sharper and sharper.

A sharp spot in a diffraction picture means you have long-range order. Ten spots in a circle means you have long-range order and *five-fold* rotational symmetry.

It turns out that people knew that this could happen, but it is impossible to have five-fold rotational symmetry and long-range periodic order, but you can have *five-fold* rotational symmetry and another type of long-range order, which is called quasi-periodic. Mathematicians get mad if you say quasi-periodic, you really should say "quasi-crystalline," and I will try to remember to do that.

[Transparency 20]

Here are some electron micrographs of this aluminum alloy and you can see here that there are 5 facets -- this is macroscopic size, a couple hundred microns, which is really macroscopic size relative to atomic scale. You can see these things are growing with *five-fold* rotational symmetry.

[Transparency 21]

Here is another picture. This aluminum alloy was made at the Alcoa Company. Apparently they make it especially just for the Chrysler Corporation. (Laughter)

[Transparency 22]

In one dimension, it is fairly easy to understand quasi-periodicity. Suppose I take a line and I put dots on it with spacing lattice constant a . If I do a diffraction measurement or just take the Fourier transform of this, then I will get a sharp line at π/a , a sharp line indicating that I have long-range order. Throughout this whole system there is just the lattice constant a .

On the other hand, suppose I take a line and put dots on it randomly, so there is no characteristic length. In that case, I will get a broad spectrum when I take a Fourier transform.

Now suppose I take another line and put dots on it periodically with a lattice constant b and then I superimpose those 2 lines. What I will get is a line with dots on it that may not look all that different from the line with the random dots on it.

If these 2 lattice constants are commensurate, that is, if their ratio is equal to a rational number (that is, it is equal to a ratio of 2 integers), then there will be some pattern here that will repeat periodically. That ratio of integers, if it is a ratio of 2 big integers, then the unit cell might

be large, the pattern might have a large number of dots in it, but that pattern will repeat periodically, and you can use Bloch's theorem and other theorems of group theory to treat that system.

On the other hand, if these 2 lattice constants are incommensurate, that is, their ratio is equal to an irrational number, then there will be no pattern in here that repeats periodically. However, if I take a Fourier transform, then I get 2 sharp lines -- this was just the linear superposition of these 2 periodic systems -- at π/a and π/b . The sharp lines indicate you have long-range order, but because these spatial frequencies are incommensurate, you do not have any periodic order.

That is a quasi-periodic system. Quasi-crystalline systems are a little more sophisticated. Certainly in higher dimensions, quasi-crystalline systems are quite special.

To explain how you get a quasi-crystalline system in higher dimensions, let me first show you the way you get a quasi-crystalline system in one dimension like this quasi-periodic system here.

To get a quasi-crystal you start off with a periodic system in a higher dimension, so here I will start off with a square array of dots in 2 dimensions. Then you intersect that higher dimensional system with a lower dimensional surface. That surface intersects with direction cosines that are irrational numbers.

For example, I will take this periodic square array of dots and intersect it with a line. The cosine of this angle, or the tangent of this angle, will be an irrational number. Then what you do is you define a window and for the dots in the periodic system within that window, you project it onto the lower dimensional surface. Here I project these dots down onto the line.

Those dots, then, will form a quasi-crystalline sequence. Because you are projecting dots from a periodic higher dimensional system, the dots and the line will have long-range order, but because the direction cosines or the tangent of this angle is an irrational number, there will be no pattern here that repeats periodically. This is how you make a quasi-crystal.

For this 1-dimensional example, there is a special case, where the tangent of this angle is equal to the Golden mean, which is the square of $5+1/2$. It is the Golden mean, it is also called the divine ratio, and it is also called the most irrational number. I think it is that because if you did a continued fraction approximation for this irrational number, I think it is the continued fraction that converges the slowest.

If the tangent of this angle is equal to that, then the sequence of dots that you get on the line correspond to a sequence of Fibonacci numbers. As you probably learned in high school, Fibonacci numbers have all sorts of amazing properties. By the same token, this 1-dimensional quasi-crystal, based on the Golden mean, also has really amazing properties.

To make a 3-dimensional quasi-crystal, you start off with a 6-dimensional periodic lattice and you intersect it with a 3-dimensional surface, and I am sure you will have no trouble at all imagining that. (Laughter)

You project dots onto that 3-dimensional surface and you will get a 3-dimensional quasi-crystal.

To make a 2-dimensional quasi-crystal, you start off with a 5-dimensional periodic system and intersect it with a plane and project dots and you will get a 2-dimensional quasi-crystal. Two-dimensional quasi-crystals are also known as Penrose tiles.

[Transparency 23]

These date back to quite a bit before the discovery of the aluminum alloy quasi-crystals. Roger Penrose was the first to play with these. There were a number of articles in *Scientific American* written by Martin Garner (some of you remember him). This is a standard, one of the common Penrose tile patterns. This one you can see has *five-fold* rotational symmetry.

To make a crystal, you take one unit cell and repeat it to fill all space. To make a quasi-crystal, you are allowed to use more than one unit cell, so here there are 2 unit cells, there is a fat rhombus and a skinny rhombus, and the ratio of the areas of those rhombuses is the Golden mean, so there is an irrational number involved in the size scales here.

By the way, I should mention some other references. There was a cover story in *The American Scientist* on these things and, also, there is -- or at least there used to be -- a company in Texas that made ceramic tiles like this that you could use to tile your bathroom floor.

In order to tile the plane without leaving any gaps, you have to put these tiles down -- you have to follow a set of rules. If you do not follow those rules, you may get to a point where you cannot fit a tile without leaving a gap. Because of these rules, you have long-range order, but because of this irrational number that is involved in the size of the 2 unit cells, there is no pattern here that repeats periodically.

MR. SMITH: Seeing that makes me think of sidewalk cracks and I remember, depending on the speed and propagation, if it is a fast break, you get, I believe, 4 lines coming out and if it is a slow break, you get 3, or it is the other way around, and so --

DR. MAYNARD: These things do occur in nature when you do this rapid quenching of these aluminum alloys. It was a mystery as to how nature figured out how to make these quasi-crystals, because the way I described it, this mathematical technique of starting off in 6 dimensions with a periodic lattice and then projecting onto a 3-dimensional surface -- I mean, crystals do not grow that way.

Someone did figure out how you could use a local set of rules, and this is like how a crack is going to branch, but there is a local set of rules so that nature can grow these quasi-crystals just using these local rules and letting it grow out. That is really pretty amazing. It took them quite a while to discover that. But yes, you can have local rules that will give you this *five-fold* kind of pattern.

DR. HOFER: Is there anything special, in the metallurgy case, the aluminum, about how quenching does this? I know if you quench some certain alloys really fast you get metallic glasses that are amorphous and then you go from that to polycrystalline. What causes the quasi-crystalline to come out?

DR. MAYNARD: I do not know. I am not sure anybody really knows the atomic details about how it gets figured out, how it does it, but this does have closely related cubic crystals. If you change the stoichiometry, the amounts of the constituents, a little bit, there is a region in the phase diagram where it forms a quasi-crystal.

MR. WAKELAND: Does anyone know whether there are any more of these or not?

DR. MAYNARD: I think there are probably an infinite number of patterns that you can make with 7-fold, 13-fold, and so on. In fact, I think 10-fold is not allowed in classical crystals. There is a whole number of ways you can do this, and rules, and so on, about forming them.

This thing has no periodic symmetry, so Bloch's theorem does not apply, but there are some other theorems that apply to these things. There is a theorem that says if I take these 2 unit cells and draw lines on them in some way, then when I put these tiles down according to the rules, the lines that are on the tiles will also make a Penrose tile pattern, but it will be scaled down, and that is called inflation or deflation symmetry.

An even more amazing theorem says that if you pick out some pattern with some nominal diameter, then within 2 diameters, and usually within 1 diameter, you will find that pattern again, but it will not recur periodically. That is pretty amazing. First of all, how did they discover that? Somebody must have stared at quasi-crystals for a long time. Then he had to prove a theorem that showed that.

In this one I have this global *five-fold* symmetry but, in general, a quasi-crystal in real space does not have an axis of *five-fold* symmetry; it is only in the Fourier transform space. For a more arbitrary quasi-crystal, finding a pattern and then within 1 or 2 diameters finding that pattern again I think is pretty amazing. That is a neat theorem. I think that theorem has a lot of consequences as to the properties of the quasi-crystals. It is called Conway's theorem.

[Transparency 24]

Now there are 3 fundamental types of solids: There is periodic, or crystalline solids; there is disordered, or amorphous solids; and now there is quasi-crystalline solids.

A fundamental question is suppose you have a wave equation with a quasi-crystalline potential field. How does this symmetry show up in the eigenvalues and the eigenfunctions? It is a very fundamental question to ask.

In 1 dimension you can prove rigorous theorems. This has to do with the Golden mean and the Fibonacci numbers. In 1 dimension you can prove rigorous theorems.

[Transparency 25]

The eigenvalue spectrum forms a Cantor set, which of course is related to M being a manifold. There are bands and gaps everywhere and the ratio of the amounts of bands and gaps is the Golden mean, again. The eigenfunctions can be extended eigenfunctions, they can be localized eigenfunctions, or they can be so-called critical eigenfunctions, which is neither of those, but that is only in 1 dimension that you can prove rigorous theorems.

[Transparency 26 – unavailable at time of printing]

In 2 and 3 dimensions no one has been able to prove any rigorous theorem. It is not periodic, so you cannot use Bloch's theorem, but it is not random, it has long-range order, so you cannot use statistical theorems, either. No one has been able to prove a quasi-Bloch's theorem for these quasi-crystals.

[Transparency 27]

With no rigorous theorems, we thought we would do an experimental measurement to find the properties of a 2-dimensional quasi-crystal. What we wanted to do was to once again use a tight-binding approach, where you have a local oscillator that has a definite frequency or frequencies and then you put the local oscillators in this Penrose tile pattern and then you couple them together.

When you couple them together, the eigenfrequency of the local oscillator will broaden out into a whole spectrum of frequencies. Then you want to look at that spectrum of frequencies and see if there is any structure that reflects the quasi-crystalline symmetry.

Again, in these systems you want to use something with low damping, because if there is damping, you will not get this long-range phase coherence, the interference of the waves that produce the structure that you are looking for. We wanted a high-Q local oscillator.

If you want a high-Q local oscillator, it is hard to beat the tuning fork. I can get this started and it will ring for a couple of minutes. I am holding it at the stem here, and that is a lossy support, but the 2 tines work against one another, so there is not much motion here, so not much energy leaks out at the stem.

Also, you cannot hear it, and that is a good thing, because that means it is not losing energy by radiating sound out into the room. This is a high-Q local oscillator (its Q is about 10,000 or up).

After buying a spool of music wire, we went back to the local music store and bought 300 tuning forks. (Laughter)

The guy at the store was very interested in what we were doing with this. Once we did the experiment, made the measurements, and got it published, I went back and showed him the paper and he was quite interested.

Here is our local oscillator. The frequency is 440 Hz. They are made to be identical. To couple them together we used an arc of steel wire between the tines and spot-welded it to the tines, so when the tines of neighboring tuning forks moved, it bent the steel wire and that provided the coupling between the tuning forks.

PARTICIPANT: Are those all lined up linearly?

DR. MAYNARD: No, this is poetic license the way they are lined up here in the drawing. The way they were actually aligned is here is a rhombus from the Penrose tile pattern. The tuning fork was aligned so that the 2 tines were across the short diagonal. Then the 4 sides of the rhombus identify 4 nearest neighbors. There were coupling wires to the nearest neighbors like that.

Because of the different sized tiles, the 2 different rhombuses, it is a little more complicated but, still -- this is actually like a diatomic molecule in a solid. It is a complicated local oscillator, but the fact that it has the Penrose tile symmetry is inescapable. The coupling wires, when you calculate the distances to any nearest neighbor, were only 4 different lengths that were involved.

These tuning forks are epoxied into a 3/4-inch aluminum plate with about a meter and a half on a side. We thought at first we might get coupling just through the stems and the base but that was really small. We needed to get strong coupling so that the 440 Hz would broaden out wide enough to where we could measure the individual resonances.

PARTICIPANT: How do you control how strongly coupled they are to get to the physical--

DR. MAYNARD: We made the coupling wires on a jig. There are 4 different lengths, so we had 4 different jigs, so each wire was made as similar as possible, and then we were just very careful when we spot-welded them to the tuning forks.

PARTICIPANT: Your coupling between the tines, you did not actually make, say, little puzzle pieces and fit them together, you just had a drawing?

DR. MAYNARD: As I say, this thing makes it pretty clear as to where the wire goes and how it goes. As I say, there are only 4 different lengths that turn up because of the Penrose tile pattern. It was tedious but straightforward.

To make a measurement, we put an electromagnet next to one of the tines and put an A.C. current through the magnet. To measure the response we used electric guitar pickups -- a lot of you graduate students will appreciate that. We went back to the music store, bought half-a-dozen electric guitar pickups, and put them next to the tines of different tuning forks in here. Of course, they will give you a nice signal from the vibrating tuning fork.

We would sweep the frequency and measure the response of the system and we would get resonances at the eigenfrequencies for the coupled system of tuning forks in the Penrose tile pattern. The results are shown down here. This gives you the frequencies.

Here is the 440 Hz. There are bands and gaps here. In these regions, here there are no eigenfrequencies. The 440 Hz, the original local oscillator frequency, winds up in a gap. That is pretty common in solid-state physics. If you have a solid with a unit cell that has 2 different atoms, like sodium chloride, then you can have phonon modes where the 2 atoms move together (those are called acoustic modes) and you can have modes where they move opposite (optic modes). The same thing can happen with the tines of the tuning forks, you can have motion like this or motion like this.

This would be like an acoustic branch in solid-state physics and this would be like an optic branch, with the local oscillator frequency in a gap.

We saw very definite gaps in the spectrum of frequencies. You would be sweeping the frequency and you would get a resonance and then another resonance, a resonance, and then nothing. You can turn up the drive, you can turn up the gain on the amplifiers, there was nothing going on at those frequencies.

With 300 tines in the array there should be 300 eigenfrequencies. With the tuning forks and the high Q, if you look at this, this is only a little over 100 Hz wide and there are 300 lines in there and we could resolve every line because the Q was so high. The width of each line was about a 10th of a Hz.

DR. HARGROVE: Were all of these lines taken off of one fork?

DR. MAYNARD: No, we measured about 20 different spots.

DR. HARGROVE: I ask it, because I did not think that was clear.

DR. MAYNARD: Yes, we had 6 electric guitar pickups but we moved them around to make sure that they were not sitting at nodes. If you are sitting at a node, then you miss the response.

DR. HOFER: You just said that the number of modes you count is equal to the number of tines, not the number of tuning forks?

DR. MAYNARD: Yes. You have 300 Newton's laws, 300 simultaneous equations, and 300 eigenfrequencies. We could actually see them all.

These gaps were real, so you have these bands and gaps. If you take the ratio of the width of any of those bands and gaps, you get the Golden mean, or the Golden mean to some power. That was something that was not predicted by the theory, because there is no theory for 2 and 3 dimensions.

People did computer simulations of quasi-crystals, Penrose tiles, and they did not see this effect. There is a lesson there -- it is a lesson for solid-state physicists, not for us acousticians so much. In the theory, they used a thing called a hopping Hamiltonian. A hopping Hamiltonian is basically a matrix and in the matrix for this thing, they would have either a 1 or a 0, depending

on whether or not there was a nearest neighbor in the Penrose tile pattern. That has the topology of a Penrose tile pattern but it has nothing to do with the wave equation.

[Transparency 28]

I think the reason you get these bands and gaps is that in our system, with these coupling wires, the coupling wires have a finite mass, so there is a finite speed of sound in our system. As you vary the frequency, you are varying a wavelength. By varying the frequency and varying the wavelength, you are probing all the length scales in that Penrose tile pattern.

The Golden mean is in all those length scales, that is how it is constructed. In that sense, it is not surprising to see the Golden mean appearing in the frequencies, but in this hopping Hamiltonian that setup is equivalent to having an infinite speed for the wave. In that case, you are not probing any of the length scales and you do not see the Golden mean.

With a finite speed for the wave, at some frequency you might be able to fit -- see, there are some patterns, like this pattern -- an integral number of wavelengths in that pattern. It will like that and lock up there. By Conway's theorem, this pattern recurs. I think that is why you are guaranteed to see these bands and gaps but, as I say, it was not predicted by the theory. We had to spend the 6 months in the lab to measure it first.

MR. CURRA: How do you decide which one is a band and which one is a gap? I am looking at this and at the previous slides and you can pick a lot more bands than gaps in-between the --

DR. MAYNARD: That is right. For a finite system, you are always going to have discrete frequencies. As you say, what do you decide is a band or a gap. We just took the biggest spacings. If you take the limit, as the thing gets infinite, then it becomes continuous, but for the Penrose tile pattern it may become continuous and the gaps may actually shrink to 0. In the limit there is still a ratio that gives you the Golden mean. That is what happens in 1 dimension. This is a lot like the result in 1 dimension.

MR. CURRA: The gap becomes bigger independently of what the bands do. Is that a coincidence?

DR. MAYNARD: I have no idea. There is no theory that gives you any guidance here.

DR. HARGROVE: Is there any tiling that gives you a fully tiled surface that has a different ratio from the Golden mean and, if there were, would you expect to find that irrational thing to show up in the spectrum?

DR. MAYNARD: Yes, almost certainly. In fact, if you do it properly -- you can take masses and springs and couple them together in a Penrose tile pattern and you can write down all the Newton's laws, and so on, using the transmission line formula to go from one on to the next.

You will have trig functions and the argument of the trig function will be ω/C times the length. When you build these things, whatever rational number you use, those lengths are going to have that irrational number in them. When you take the determinant of this matrix to get the eigenfrequencies, that irrational number should show up in the frequencies.

That is not a proof but it makes it plausible.

DR. HAMILTON: For long wavelengths, when you have a sound wave going through a quasi-crystal, do you model that as an isotropic solid? What do the stiffness constants look like?

DR. MAYNARD: Oh, yes. In fact, that was Phil Spoor's Ph.D. thesis. He got a real aluminum alloy quasi-crystal and used resonant ultrasound to measure its elastic constants.

DR. HAMILTON: For example, for crystals you have only, say, 21 independent --

DR. MAYNARD: The theory for a 3-dimensional quasi-crystal -- in fact, you can use group theory, group theory predicts that they will be elastically isotropic, 2 moduli. In crystals, the highest symmetry you can have is cubic, but for a cubic crystal it is not elastically isotropic. You need 3 elastic constants for a cubic crystal. Quasi-crystals are special in that they need only 2.

DR. HAMILTON: Can you get down to a length scale at high frequencies where all of a sudden it takes on a crystalline behavior, I guess?

DR. MAYNARD: I do not believe so. As I say, it is group theory; it is just rotational symmetries that do this.

DR. MOLDOVER: Is that the way it turns out?

DR. MAYNARD: It turns out that the quasi-crystal, these aluminum alloys, if you change the stoichiometry just a little bit, there is a very similar cubic system. Because the quasi-crystal was rigorously isotropic, the cubic system is anisotropic but not very much. You have to make very precise measurements in order to see the difference between the isotropic quasi-crystal and a cubic approximate.

Phil Spoor had to make very accurate measurements just to see it, but he did.

Back to the tuning fork Penrose tile. These are the eigenvalues. Remember that for any system you need to measure eigenvalues and eigenfunctions. To measure the eigenfunctions we took the tuning fork quasi-crystal and put it up on the wall of the lab.

[Transparency 29]

On the end of the tine, we put a little tiny mirror. Even when you are driving this thing at a resonance you can feel, the tines vibrate, just as with a regular tuning fork you can feel the vibration. You cannot see anything; there is not much there to measure. We put a mirror on the end of the tine, shone a laser on it, the laser beam would bounce over to the other side of the room and, as the tine would vibrate, you had this optical lever arm, so even though you could not see any displacement here, you could get a couple of centimeters displacement on the opposite wall of the lab with that large optical lever arm.

We put a mirror on every tine and we had a scanning system that would scan the whole thing. We would darken the room, open a camera so that it was looking on the opposite wall and then we would scan the whole system. We would get streaks over here that would show you the displacement of each tine.

The problem is that each mirror has to be adjusted so that the reflected laser beam goes to the corresponding spot on the opposite wall. There was a poor graduate student who was up there for months and months aligning 300 little mirrors, but he did it.

[Transparency 30]

Here are some of the eigenfunctions. The trouble is that when you have a system that has some perturbations in it -- again, you can use perturbation theory to show that the eigenvalues are pretty robust, the perturbations do not shift the eigenvalues very much, but perturbations do affect the eigenfunctions.

Those little spot welds were not perfect and they did not affect the eigenvalues very much, but they did affect the eigenfunctions, so the eigenfunctions we measured were not great. We could see some qualitative effects. Here is one that has the amplitude distributed sort of

uniformly. There were others where you could see some *five-fold* symmetry, and others that had large amplitudes localized. We think these might be examples of localized states.

It was only a finite system, so we could not really draw too many conclusions from the eigenfunctions, but the eigenvalues were a good measurement.

[Transparency 31]

This was a pretty popular experiment and got written up in *The New York Times*, with a picture of the tuning forks. It was a lot of fun to have that happen.

The next topic is going to be some nonlinear effects, so before I do that, why don't we take another break.

DR. MAYNARD: I hold in my hands the last set of viewgraphs for the summer school. (Laughter)

[Transparency 32]

I have talked about disordered systems. An interesting question comes up: What if you have a system that is both disordered and nonlinear? For the disordered systems there has been a lot of research, there is the Anderson localization and these other effects, lots of *Physical Review Letters* articles, there has been a Nobel Prize, and 2 Buckley Prizes involved here.

At about the same time, but independently, there has been a lot of research on nonlinear systems. Professor Lauterborn told us a lot about these systems. There are solitons and how to solve nonlinear equations using inverse scattering theory. There is also chaos that Professor Lauterborn talked about. Of course, with chaos there have been a lot of books written about chaos and solitons. There has also been some great humor about physicists and chaos.

[Transparency 33]

Here is one. I think it was Jim Mehle who sent me this. It says, "What led you to the mathematics of chaos, Dr. Maynard?" They got my name and they also got a pretty good likeness, I think. (Laughter) Although those of you who know me probably know that my office is a little neater than that.

[Transparency 34]

There has been a lot of research for waves in disordered fields, nonlinear waves. The question is what if you have a system that is both disordered and nonlinear. There should be a lot of interesting physics in the overlap of those 2 fields.

Let's take a look at the research that has taken place for systems that are both disordered and nonlinear. It turns out that there has not been much research. The reason is that this is a very difficult field. Nonlinear physics is difficult enough by itself, as are disordered systems, so the study of systems that are both disordered and nonlinear is very difficult and it is also a quite recent field.

The first meeting on systems that are disordered and nonlinear took place in 1989. Since then I think there have been only another couple of meetings. At your nonlinear meeting, are you going to have anything about disorder there as well?

DR. LAUTERBORN: We do not yet know.

DR. MAYNARD: It is a very new field. A lot of the people who work in this field are hard-core mathematicians.

For systems that are both disordered and nonlinear, there is a very fundamental question and that is, does nonlinearity weaken or destroy Anderson localization? There have been only about 8 groups that have attacked this question.

By the way, I should mention that at these meetings on disorder and nonlinearity all of the papers have been theory papers, so there is lots of good opportunity for experiments here.

For this question of does nonlinearity weaken Anderson localization, there are 8 theoretical groups that have addressed this problem. As you can see, some of the groups say yes, nonlinearity will weaken Anderson localizations. Some say no, it will not. Some vote both ways. (Laughter)

It turns out there is not a big controversy here; these people really are not cheating. It is a nonlinear system and some nonlinear systems do not have unique answers or they have answers that depend very precisely on how you ask the question. That is the case here.

[Transparency 35]

This shows a couple of ways you can ask the question of disordered systems. There are 2 ways you can analyze a system. You can do a normal mode analysis. In a normal mode analysis you drive it at 1 frequency, $\cos \omega T$, and then at the end of the system you look at what comes out, so you have a transmission coefficient. You do that as a function of ω .

You look at all frequencies and if you are at a frequency that corresponds to an eigenfrequency of the system, you will get a peak in the response, in the transmission. As a function of frequency, you will get a transmission spectrum. You can analyze the system by driving it $\cos \omega T$, sweeping ω and getting a transmission spectrum.

On the other hand, you can take a system and send it a sharp pulse and that sharp pulse will rattle around in the system and then, at the exit, at some point, you will get some displacement as a function of time and you can just look at that.

For a linear system these 2 things are Fourier transforms of one another, because the drives here are just Fourier transforms of one another. At delta function peak the Fourier transform is the whole spectrum and vice versa.

PARTICIPANT: How do you say the name on the pulse?

DR. MAYNARD: I am not sure. Goupillaud? Is that close? [No response]

It is a special disordered system and a lot of people studied this system because it was easy to do the calculations but it turns out it has pathological results. If you send in a pulse, you get something like this, whereas in a real disordered system, if you send in a pulse, you get this. So watch out for papers that do this.

Each section here has a different speed of sound, but the distance of each section is adjusted so that the transit time is the same. In some sense, it is disordered, because the speed of sound is disordered in each layer, but they adjust it so the transit time is the same, and that has some really special effects. A real disordered system looks like this.

For a linear system, doing the normal mode and the pulse analysis are just a Fourier transform of each other, but in a nonlinear system that is no longer true, so a normal mode analysis can give you a different result from a pulse analysis.

We wanted to study this with our acoustic analog systems. First, we decided to take a look at the normal mode, do a normal mode analysis of our string with the masses on it.

[Transparency 36]

Let's take a look at the nonlinearity of a string. If I take a string and then stretch it to some tension T_0 and clamp it here, transverse displacements of the string obey the wave equation and from right here you see that the speed of the wave is the square root of the tension over the mass-per-unit length, but now suppose I take the string and I give it a finite transverse displacement like this.

In that case, the arc length of the string has increased, so the string has been stretched and the tension in the string goes up. If you calculate the actual tension in the string with its displacement -- there is the tension without the displacement. This thing here is essentially the integral that gives you the total arc length of the string and this would give you the Young's modulus for the string.

DR. HARGROVE: Is that at steady state?

DR. MAYNARD: This would be if you did a static transverse displacement, for example.

DR. HARGROVE: If you make a dynamic thing, it has --

DR. MAYNARD: Yes, right.

An interesting thing here is that this is a global nonlinearity, because even if I have a localized pulse here, at least in the static case, the tension changes in the entire length of the string.

As Logan mentioned, suppose you have a dynamic system, where this is oscillating. This change in the tension gets propagated along the string -- remember, in our experiment this is a steel wire -- at the speed of longitudinal waves in the steel. The speed of sound in the steel is a lot higher than the speed of these transverse waves, so a local displacement, even for the dynamics of the transverse waves, the change in the tension is distributed almost instantaneously over the whole wire. This looks like a global nonlinearity for a localized wave.

What we wanted to see happening was the following. Suppose I have a disordered system and I have Anderson localization at some site like this. If I drive it at finite amplitude, then it is going to modulate the tension in the wire at the frequency of this state.

Robert talked about parametrically driving systems. If you take a string and just pull on it like this, you might say, well, you are just pulling it straight, so nothing happens, but it is unstable to the excitation of transverse vibrations.

If there were some little thing that made it displace this way, when I pulled on it, it would come down, but then I let it in and it falls through. I pull on it again and it comes up. You can see how you can parametrically excite a transverse vibration. Of course, if you do it at twice the frequency, you drive the resonance of that transverse vibration.

Here is what we wanted to see happen. You have Anderson localization at one site and it has a finite displacement, so it parametrically modulates the tension in the string, but then the modulation of the tension can parametrically excite another Anderson localization state at a different site.

What we would have, then, is an Anderson localized state but the nonlinearity can cause it to hop to another Anderson localized state. This state here runs at some frequency, so it modulates the tension at twice that frequency, but this other eigenstate, being another eigenstate, has a different eigenfrequency. You might think that this modulation of the tension is at the wrong frequency.

Remember that Robert also showed us a stiffening system where, if you go to finite amplitudes, a tuning curve will bend over. Now suppose we have 2 nearby resonances. Now when they bend over you can have either one of these states being driven at the same frequency. One of them will be at a higher amplitude at that frequency but the other one can oscillate at a lower amplitude at the same frequency. It could be that you could have resonant parametric excitation of a different Anderson localized state.

People did normal mode analysis for disorder, for nonlinearity and disorder. The prediction was that the nonlinearity for the normal mode analysis would not destroy the Anderson localization, but this parametric hopping was something they could not treat in the theory, so we thought we would do an experiment to see if we could see nonlinear-assisted hopping between these localized states.

[Transparency 37]

I will just jump to the results. It was basically the same experiment with the masses on the wire and to see nonlinear effects you literally turn up the volume.

Here is a transmission spectrum for the normal mode analysis at a relatively low drive amplitude. For the disordered system you get these various eigenfrequencies. As you increase the amplitude -- we took this transmission spectrum and divided by the drive amplitude, so they are all normalized.

If it were a perfectly linear system, each one of these curves should be the same, but you can see that they change. In fact, if you follow this one here, it is going up in amplitude quite a bit. Then it drops down. That one goes up but then it drops down. It is like those tuning curves. You go up the tuning curve but then it drops off the end. The same sort of thing is happening, I think. This is in a disordered system.

Now, even though it looks as if some of the eigenfunctions are getting larger amplitudes here, if you integrate to find the total transmission through the system, it turns out that if you normalize to the drive the total transmission is dropping. If we had this nonlinear-assisted hopping, the transmission would increase. If we have a state localized over here far from the receiver we get small transmission, but if we get the nonlinear-assisted hopping, then we will get larger amplitude at the receiver.

It is not really what is happening here and, in fact, you do not get that effect. Yes?

DR. MOLDOVER: The drops in amplitude, were they reversible?

DR. MAYNARD: I cannot remember.

DR. MOLDOVER: Because the bending curve -- you would expect hysteresis.

DR. MAYNARD: I will show you some data that show that for a multi-resonant system.

[Transparency 38]

Let me show you some actual eigenstates. Here is an eigenfunction (it is one of those that was increasing) at a relatively low drive. If you increase the drive, you can see that some of those local oscillators get much higher amplitudes, but if you look at this localization length, the localization length here has not really increased.

What happens is as you increase the amplitude, instead of a localized state parametrically exciting another state, it actually parametrically excites itself and, if anything, it gets more localized. We did show that for the normal modes; going nonlinear does not destroy the Anderson localization, it still remains.

Pulses are a much more interesting system and it is for pulses that you can get either answer, depending on exactly how you do the measurement.

[Transparency 39]

Here is the theory for a nonlinear pulse in a disordered system but, first of all, let me go over what you can do with a pulse in a linear system. You can imagine sending in a pulse and having it rattle around in here, but that is a difficult calculation to do.

For a linear system, the way to do the calculation is you first do the normal mode analysis or, in other words, you find the transmission spectrum. The way you do that in the normal mode analysis is you imagine sending in a $\cos \omega T$, you have a reflected wave and a transmitted wave.

To do the analysis, you assume that the transmission is 1, so that gives you a unique condition here, and then you work your way back down the system. As you go past each scatterer you multiply by a 2×2 matrix. Why a 2×2 matrix? The wave equation is second order. Between each one of these sections you have a linear combination of 2 linearly independent solutions and as you go across each scatterer you are multiplying by a 2×2 matrix.

You multiply by these 2×2 matrices from one end of the system to the other end and then you can calculate the incoming and reflected and you normalize to the incoming. This value then becomes the transmission coefficient and this value becomes the reflection coefficient. If you are doing a pulse, you take a Fourier transform for a linear system.

The interesting thing is that even though you have a pulse here, in order to understand what that pulse does you have to multiply 2×2 matrices from one end of the system to the other. By Furstenberg's theorem, you get Anderson localization if it is a disordered system. The title of Furstenberg's paper is something about the product of random matrices.

For a linear system, a pulse, this analysis guarantees that you are going to see effects of Anderson localization and for a pulse what it means is that when the pulse goes in, its amplitude exponentially decays with the characteristic Anderson localization length. That is for a linear system.

If the system is nonlinear, if you have a nonlinear pulse, it turns out a nonlinear pulse has extra degrees of freedom. For example, if you do the analysis with the inverse scattering, the extra degrees of freedom or the number of eigenfunctions that you get in the inverse scattering problem is the number of eigenvalues that go into making the nonlinear pulse.

With the nonlinear pulse, you have extra degrees of freedom, so you can adjust those extra degrees of freedom so that you have to satisfy only conditions locally. For the linear pulse you have to satisfy conditions throughout the entire system. In other words, with a nonlinear pulse can you define a second characteristic length and for a soliton, it is basically just the width of the soliton.

Now when you have disorder and a nonlinear pulse there are 2 characteristic lengths; there is this nonlinear length and the Anderson localization length. What happens depends on the relative sizes of those 2 lengths.

If you have a strong soliton or a strong nonlinear pulse so that the nonlinear length is less than the Anderson localization length, then the strong soliton does not really sample enough of the disordered system to realize that it should be Anderson localized. A strong soliton will just propagate through a disordered system.

On the other hand, if you have a weak soliton, so that this nonlinear length is much larger than the Anderson localization length, then it knows all about Anderson localization, it sampled a large section of the disorder, so a weak soliton suffers the exponential decay from Anderson localization.

This is a plot of the log of the transmission as a function of the distance traveled. If you have localization, you get this exponential decay of the amplitude of the pulse.

The interesting case is when the nonlinear length is on the same order as the Anderson localization length. In that case, the theoreticians predict that this pulse will go along for a while with some decay, but then there will be a break and they will start to exponentially decay. This is the behavior we wanted to try to see in an experiment.

[Transparency 40]

For this we needed to have a local nonlinearity, so the nonlinear waves on the string were not going to work here. If you want a local nonlinearity, it is convenient to use surface waves, because for a surface wave the speed of the wave is the squared of the depth times the derivative of the potential at the surface. For gravity that is G times D , so these are shallow surface waves.

The important point is that the speed of the wave depends on the depth. If you have a surface wave present, in addition to the equilibrium depth you have the depth under the wave itself, so the speed should depend on the depth and the amplitude of the wave itself, the surface wave.

If you plug this into the wave equation, do a Taylor series expansion, you can clearly see you are going to get a nonlinear wave equation; so surface waves are intrinsically nonlinear.

We also want this system, in addition to being nonlinear, to be disordered. In order to see the effects of the disorder, say, for low amplitudes, you need the long-range phase coherence; in other words, you need low attenuation of the wave.

We decided to use surface waves on superfluid helium, which has no viscosity, and such waves are called third sound. You have a substrate with a thin film of the helium, so there is some normal fluid helium present but it is clamped by its viscosity to the substrate and only the superfluid moves.

The restoring force for such a thin superfluid helium film is not gravity, it is a van der Waals force, so this potential here goes like some coefficient over distance cubed. The third sound is well understood theoretically. We could make finite-amplitude surface waves on superfluid helium. There is relatively low damping, because it is superfluid.

For the disordered potential field we took this substrate and cut notches in it. We made one substrate that had a periodic array of notches; we always wanted to be able to see Bloch wave states, because they are theoretically completely understood. We had another substrate that had a disordered array of notches down the length.

We had transducers that would launch a plane wave of third sound and a receiver to detect the amplitude of the third sound.

[Transparency 41]

Here are some results for the periodic system. You send in a pulse and when it gets down to the end after all the periodic scatterers you see a main received pulse and then some things that look like reflections. You might think this is just the pulse rattling back and forth in the periodic scatterers like echoes coming out.

If you think about it enough, you can see that that cannot be happening, because the first pulse here would be decayed by a transmission coefficient for one scatterer to the n th power, so it could not really look like this. If you do a proper analysis of a pulse in a periodic array of scatterers, you will get something like this.

In fact, for low amplitudes, if you take the Fourier transform of this, you see the band structure. We have some better data than this but this is where we had finite amplitudes.

[Transparency 42]

Here is a band here, and this is a gap, a small gap here, another band, small gap, and so on. This is a nice result for a periodic system for a low-amplitude pulse.

Then we increased the amplitude of the pulse and you can see this structure goes away. If you take the Fourier transform, you can see that the band structure goes away. What is happening here, as we make the pulse more and more nonlinear, that nonlinear length is getting smaller and smaller, so it is sampling only a small fraction of the periodic system. It does not really care about Bloch's theorem; it is seeing only a few scatterers. That was a good test that things should work right.

[Transparency 43]

Here are some more results for the nonlinear pulse. This is at different drive levels and then it is normalized to the drive levels, so they have about the same amplitude in their received signal. You can see that the transit time of the pulse is changing. The speed of a soliton depends on the amplitude. You can see that here. This is definitely an indication that we were seeing nonlinear effects.

[Transparency 44]

We did measurements on the disordered system. The difficulty there is that when theoreticians study disordered systems, each disordered system has a unique answer. What the theoreticians do is they ensemble-average; they assume that they have a large number of systems with the same distribution function for the disorder but not the same realization of the disorder. They ensemble-average and they can pick one number.

In the experiment, we have only one realization of disorder, so we had to do a number of measurements and then do experimental ensemble-averaging, which is quite difficult. It was very tedious.

[Transparency 45]

We finally got some results and they are shown here. These results are more recent than the ones in the handouts. The points on the handouts are not quite as good as this. We were able to do more ensemble-averaging. In our system we could verify that the Anderson localization length was about the same order as this nonlinear length. You can see that the data do pretty much follow the theoretically predicted line, so that was nice.

[Transparency 46]

There is another experiment we did on nonlinear effects that goes back to the steel wire with the masses on it. What we wanted to see was related to what Robert Keolian showed for the nonlinear spring. If you have a spring that gets stiffer with amplitude like this, then the tuning curve bends over; with the resonant frequency at higher and higher amplitudes the spring gets stiffer and stiffer. If you sweep up in frequency, you go up to the top of the tuning curve and then it falls, you sweep down in frequency, it will jump up like that.

Suppose you have a wire with a lot of masses on it. Then, remember, you get band structure and within each band, if you have a steel wire with, say, 24 masses on it, then each band has 24 resonances in it and they can be packed close together in that band.

[Transparency 47]

What we wanted to know is suppose you have a multi-resonant system. This is a narrow band, maybe a kilohertz and the width here may be only 100 Hz or less. What happens if you have a whole bunch of resonances close together and you drive it to high amplitudes?

As you increase the amplitude, you can see each one of these resonances behaving like a simple harmonic oscillator, a nonlinear oscillator. The curves go up and they fall off; each one does that. As you go up to higher amplitudes, it gets more complicated but it is still the same thing that it goes up and drops off.

MR. POESE: If you have two states, if you approached them from the left, there would be that hysteresis?

DR. MAYNARD: Yes, in fact, we did that; you see the whole thing.

The thing that came up that was completely unexpected was right here at the beginning of this system of resonances it got very busy. This is not a function of time, this is a function of frequency. What happens here is we will drive it at one frequency and we will drive it $\cos \omega T$ and the system responds $\cos \omega T$. Then we changed the frequency slightly to do this spectrum and the amplitude changes dramatically.

If you are on a curve like this, as you change the frequency, it should just go up in amplitude until it drops off like this. There is no period doubling. We are driving at $\cos \omega T$ and it is still responding at $\cos \omega T$ and there is no second harmonic generation. All that is happening as we change the frequency is the amplitude is varying as a sensitive function of the frequency.

When we sweep down, it still happens at this end of the spectrum. This was a new effect and we wanted to try to understand that. First of all, we wondered whether it was artifact of the tension not being constant or just some perturbations, so we actually did a computer simulation.

We had a bunch of masses coupled by nonlinear springs and in the computer simulation, we saw the same sort of thing. As you go up in amplitude, at this end of the spectrum, it gets more complicated and we can actually zoom into that region there. This is the very first resonance. When it drops off at the very first resonance it gets very busy as a function of frequency. It is not chaos; it is not any period doubling. We decided it was not an artifact of the experiment; it is a real effect.

In order to understand this nonlinear system, the string with the many masses on it, we needed a formula, an analytic formula for the linear system.

[Transparency 48]

The idea was this. If you have a linear system, then the response of a periodic system with 24 masses would look like this. At any given frequency ω that has some amplitude -- this has damping in it, too, so the amplitude is set by the damping; otherwise this would be 24 delta functions. If you have damping, the amplitude of the response is some function of the frequency.

In the nonlinear system, the speed of the wave depends on the amplitude. What we wanted to do is pick an amplitude that gives you a speed of sound. You put that into this function with

the frequency and that gives you an amplitude. When you are at a condition where these amplitudes are the same that is a point where the system will run.

We needed an analytic formula for this spectrum; this is what you would measure experimentally.

[Transparency 49]

This was a periodic system of masses, so the solutions should be these Bloch functions. You should be able to get an analytic formula for Bloch waves. The solid-state textbooks make it sound as if it is really easy, they just write down the answer like that.

This is a finite system, it does not have periodic boundary conditions; it is actually clamped at the end. We needed an analytic formula for the driving point impedance for this system. Since it is periodic, you can work it out. There it is.

This is the numerator of the real part of the driving point impedance. It is actually quite complicated, so that is why I say do not take Bloch's theorem for granted, it is a lot more complicated than you think.

PARTICIPANT: Is this Mathematica?

DR. MAYNARD: No, I do the stuff by hand. [Applause]

It turns out this is not quite that complicated because you can define 5 or 6 different things and you can write it in terms of those 5 or 6 different things, but I wrote it all out so you can see what is involved. It is complicated because there is damping in the system, so the Q , for example, is imaginary.

One way you can do this is with this analytic formula. The other way is by multiplying by the 2×2 matrices, just doing it numerically. I wrote a program that checked it and it did not agree. I realized that when I typed in the computer program I typed in a sine wrong.

This now is both the matrix multiplication and the analytic formula give you the same plot.

[Transparency 50]

As I said, the amplitude is a function of the frequency and the speed of sound but the speed of sound itself is a function of the amplitude. We pick an amplitude, we calculate this, and when the amplitude is the same as the one we put in, we plot a point.

Here are the bent tuning curves for the multi-resonant system. The interesting thing is that at a sufficiently large amplitude all these resonant curves bend over top of one another. At this particular frequency, you have 9 different normal modes that can run at the same frequency; they just run at different amplitudes.

This means that if you change the frequency it might fall onto any one of these modes with a different amplitude. That is why, when you change the frequency, the amplitude seems to change a lot. Except that there is still a puzzle, because suppose you are driving up here -- you are on, say, this curve, and it goes over and drops off to any other curve. Once it falls off to a different curve it is supposed to be stable.

Once it drops off, it should have followed this curve until it dropped off, but it did not do that. It jumps around to different curves as you change the frequency.

[Transparency 51]

We think this is an example of basin crowding (I probably should have talked to you about that). Let's say we are driving it $\cos \omega T$ at one particular frequency and then the trajectory, displacement, and velocity space, the phase space, is a circle -- this black circle here.

When we change the frequency, what it is doing right before you change the frequency is the initial condition for what it does after you change the frequency. We do not control that in any way; in fact, the whole thing is stepped by the computer.

Depending on what it does when you change the frequency -- these basins are basins of attraction. The basin tells you which new trajectory it is going to go to. It may be that if you change it at a particular time it will be in a basin that takes it to a low amplitude oscillation, but very nearby there may be another basin that takes it up to a higher amplitude. It may depend very sensitively on exactly when you change the frequency. That might explain the fact.

Anyway, it is all theoretical now. We have the data and I just need to talk to some nonlinear experts to figure out to write the paper.

I think I will stop here, because we have students who are going to give talks. I will conclude by saying that what you have seen are these acoustic analog experiments of serious condensed-matter physics and with \$40.00 experiments, you can publish in *Physical Review Letters*. However, the main thing about these experiments is they are really a lot of fun to do. Thank you.

RESONANT ULTRASOUND SPECTROSCOPY AND MATERIALS PHYSICS

Albert Migliori
Los Alamos National Laboratory

DR. MIGLIORI: Thanks. I am going to talk to you today about elastic constants. I will tell you what they are, why they are interesting, a marvelous classical way of measuring them, and then how you can use what amounts to simple classical physics to extract information about the details of the microscopic processes going on in solids, including changes in bond length, electronic distribution functions, and related physics simply by making classical mechanical measurements on small pieces of material.

It sounds like a lot, and it is, and the result of that relates to this comment on the opening viewgraph here.

[Transparency 1]

By the way, this is a partial list of people who have made material contributions to resonant ultrasound. I will talk a little bit about Orson Anderson, who is really the father of all of this. I merely took his good ideas and figured out how to make them fast and cheap.

In the process of doing that, we have developed measurement techniques that are slowly becoming a very popular way of measuring elastic constants that was not there before. Because of that, there is some commercial importance to this. If you work for a government lab, somehow they seem to believe that taking the taxpayers' money and actually doing something useful with it is a good thing to do.

I think it is somewhat misguided but, in any case, pieces of this work over the years have made it into industry. There are 2 companies now manufacturing resonant ultrasound systems; both of the electronics systems came directly out of Los Alamos. Because of that, it is amazing how many administrators actually supported this work. In fact, I never even knew that some of the departments at Los Alamos had actually contributed to getting our project done. Of course, they told me they had contributed only after the fact.

Then we have the Industrial Partnership Center, which is charged with communicating Los Alamos technology to the universe. Both of our state senators got involved at one point. KOB and the other television stations made contact with resonant ultrasound spectroscopy and, in fact, that was a rather amusing incident, in which we decided to measure -- we decided that we could detect flaws in structures using resonances (not surprising).

One of the structures that we believed we could make measurements on was something large, like a highway bridge. The New Mexico Highway and Transportation Department gave me the Interstate 40 bridge over the Rio Grande in Albuquerque. This is a 175-foot span concrete bridge with two vertical 12-foot eye beams. We got a 30,000-lb hydraulic shaker that we put in the middle of this bridge and we were going to destructively test the bridge. (Laughter)

Now you laugh but, in fact, we did. The bridge was slated to be torn down and they had, meanwhile, put together new spans around this bridge, so I got this bridge and we instrumented the whole thing. We instrumented it with ultrasonic transducers -- I mean, "ultrasonic," this is 2

Hz. In fact, we used microwave displacement transducers to do this. We also instrumented it with strain gauges via the civil engineering department at Los Alamos.

We made measurements of the resonances of this bridge, which were surprisingly high Q ; that is, at 2.5 Hz the Q 's were maybe 10. This was the fun part: We took this 12-foot vertical eye beam and a cutting torch (I wanted to drop it into the river, but they would not let me do that), and we put jack stands under the bridge. They were about 6 inches below the eye beams and we cut that main girder.

When we cut that main girder, of course, the local television stations were there, and our Republican and Democratic state senators, who directly contributed to this work -- as soon as it got on television -- were there. They would not, of course, appear on camera in the same frame but, nevertheless, they were both there.

We did, indeed, cut this eye beam and the noise it made when it went was pretty wild. The bridge dropped about 2 inches and then we re-measured. I am not going to show you those data. The reason I am not going to show you those data is that in spite of the fact that all of these people believed it was a complete success, I do not believe I would drive over a bridge that was tested with resonant ultrasound (but there are other things that we can test it with).

In any case, there was lots of publicity. Publicity at least keeps me funded, in some sense, and that is never a bad thing but the real reason this is fun is that the resonances of solids are a wonderful way of accessing first-class physics.

[Transparency 2]

What I am interested in are the elastic constants or the elastic moduli of solids. This is just the spring constant K for a solid in a generalized way. We have force equals $-KX$ on a spring and in a solid I have stress equals some monstrous thing times strength. Let's pick at that for a second so you understand the language and where I am coming from.

The free energy of a solid has an entropy piece and then lots of other little pieces, one of which we will put in here, which is stress and strain. This is the analog of PdV in a simple gas or a simple compressible fluid. What we are after is understanding this piece of the free energy of solids.

MR. POESE: When you say "free energy," what do you mean by free energy?

DR. MIGLIORI: Free energy is the energy actually available, so it is not just the internal energy of the solid but it also has an entropy contribution that tells you whether that energy is accessible somehow or other.

I will show you in a very interesting and critical way toward the end of the talk how this can totally confound the metallurgists and produce very simple pictures of physics when you include it, whereas if you do not include it, things are disastrously complicated.

Anyway, we are dealing with the free energy. A key piece is the stress-strain contribution. We are also going to be, throughout this talk, dealing with Hooke's law -- this is linear, stress is proportional to strain. We will make the entire process so complicated just using linear physics that you will want to throw up your hands at even considering nonlinear stuff and I am going to stay away from that.

Because we are going to be dealing with solids, typically we have a crystal structure behind it, and the crystal structure has a symmetry. The symmetries of real objects, meaning real materials, will come into the elastic response. There are also, historically, some strange bits of

notation. I am not going to write this down, because Landau and Lipshitz did a great job of it, but the relationship between stress and strain is a fourth-rank tensor, which means that there are 4 indices on each one of these little pieces.

Because of the symmetries of real objects, meaning if the magnetic field is not present and you have not done anything too strange, this stress tensor has some symmetry properties, even if the solid you are dealing with has none. You can do pretty well with just 2 indices, so I am using the conventional pairs of indices.

I am going to have to use my hands a lot in this talk. Here I have got a piece of rubber between my hands and that rubber, let's say, is a rectangular parallelepiped, it is a box, the faces are all either parallel or perpendicular to each other, and I squeeze the box. When I do that, it bulges straight out this way and squishes straight in this way.

Anything that behaves elastically like that has a set of 0's -- here -- (I will tell you what they mean in a second). What those 0's mean is if I apply a stress T_1 on the 1, or X face (1, 2, and 3 will correspond to X, Y, and Z), then I only get strains, that is, displacements, in the 1, the 2, and the 3 directions and nothing at weird angles.

However, in real solids of low symmetries, when I do squeeze it, things can move off at strange angles, and then these 0's are filled in, so I can have as many as 21 elastic moduli for a solid, but a piece of aluminum from your Coke can is a polycrystalline mess in which the intrinsic symmetries of the crystal structure are averaged over any of the wavelengths that I use for measurements and, therefore, it is isotropic and it is homogeneous.

In that case, I will have only 2 elastic moduli, one relating to compressional-like forces and one relating to shear. We are basically working in a space where we have from 2 to 21 different things to measure and it turns out that it is also very much easier to measure things with 9 or fewer elastic moduli; that is, systems in which these 0's are in place and when I squeeze it, it does not go off at an odd angle. Most of the materials I am interested in obey that and there are a lot of mathematical simplifications from dealing with only those materials, but I want to be sure that you understand that there is a lot of stuff happening when you squeeze solids that have nice crystal structures.

[Transparency 3]

Just very briefly, if you want to do a computation with the elastic tensor, for example, let's say I want to calculate the bulk modulus of a solid. The bulk modulus of a solid is a measure of how its volume changes when I apply hydrostatic pressure.

If I have hydrostatic pressure, I have no shear forces, 0, 0, and 0, and I have the same pressure on all of the faces, the same force on all of the faces, so all I have to do is solve this set of equations, which is pretty simple, especially if I have a reasonably symmetric solid. I am solving for the strains in the X, Y, and Z directions, strain meaning the change in size and, therefore, I can compute the change in volume.

I simply multiply this matrix times this vector and I get a set of equations, and I can trivially find e_1 , e_2 , and e_3 , the 3 strains, and from that I can find the volume change, so then I have the volume change for a given pressure in terms of the elastic constants and outcomes. I will not do it for you, but that is the general prescription for using these things.

[Transparency 4]

There are very fundamental and powerful connections between the elastic stiffness and thermodynamics via the free energy. Here, again, is the Gibbs free energy differential for hydrostatic pressure, where I have written PdV . If I replace it with stress and strain, we have what I showed you on the second viewgraph.

For example, the derivative of the Gibbs free energy with respect to temperature at constant pressure is minus the entropy of the solid and the derivative with respect to stress and strain at constant temperature are the lengths of the solid. (This would be volume in the case of a gas if I did this with respect to pressure.)

I can take one more set of derivatives and I get the 3 fundamental, in the sense of the 3 measurably fundamental, thermodynamic quantities. One is the specific heat, and that is the second derivative with respect to temperature of the Gibbs free energy times a T in here.

The other is a thermal expansion coefficient -- you notice these indices are the same, ii, but, in fact, the thermal expansion coefficient in a complex solid with a very low-symmetry crystal structure can have off-diagonal pieces; that is, as it shrinks, it not only comes in but it may skew and shear.

Finally, I have the elastic moduli, which are the second derivative of the Gibbs free energy with respect to stress. You see that this is a tensor and, therefore, I have a tensor quantity that contains fundamental thermodynamic information about the solid. If I can sort this out, I can probably do a pretty good job of learning neat things.

PARTICIPANT: Stresses and strains are different in this relationship than in the earlier viewgraph?

DR. MIGLIORI: Yes, these were made at different times but, basically, I can take the derivative with respect to a strain and I get a pressure, and if I take the derivative with respect to a stress, I get a volume. Here I am taking it with respect to stress to guess a length-like or volume-like quantity, so I think that is okay. You are right, though. There will be more, lots more.

[Transparency 5]

Just to pick at the thermodynamic connection a little bit, I have some scribbles here for superconductivity in a simple type I superconductor. This is a second-order phase transition when it becomes superconducting; that is, when it becomes superconducting, absolutely nothing happens -- almost. For example, the difference in free energies between the normal and superconducting state is the volume and then some silly 8π , because theorists insist on using weird units like CGS, and then the thermodynamic critical field squared.

You see, right at T_c , right at the superconducting transition, the critical magnetic field for destroying superconductivity is 0, so at the transition the Gibbs free energy difference is 0. Of course, the critical field is a function of temperature and as I go all the way down to 0 temperature I develop the full critical magnetic field for the destruction of superconductivity.

I can walk down this path this way or this way, here taking temperature derivatives, here taking pressure derivatives, and we will use pressure in this case, so the volume difference across the phase boundary is the derivative with respect to pressure of the Gibbs free energy difference and I get the volume (this is sort of a specific volume) and I get $H_c dH_c/dP$.

I can do the same for temperature and I get T/dT d by dT of ΔS -- I got ahead of myself, excuse me, so the entropy difference here is just the derivative of the free energy difference with respect to temperature.

At this point I have first derivatives and I can take a second derivative, as we showed before, to get bulk modulus change across the phase boundary -- yes?

MR. POESE: What is the critical field?

DR. MIGLIORI: The critical magnetic field for a superconductor is the magnetic field that destroys superconductivity and returns it to its normal state. Right at the transition temperature that magnetic field is 0, because nothing has happened. As I cool, it takes more and more magnetic field to destroy superconductivity, reaching 11 at 0 temperature.

MR. WAKELAND: Can you help me a little with how we got from talking about elastic properties to -- you are mentioning superconductivity now and I am not sure exactly why.

DR. MIGLIORI: Let me continue and we will get there.

We are looking down this path now, where I am taking pressure derivatives. The first derivative got me the volume. The second derivative is a derivative of the volume change with respect to pressure, and that is going to be a stiffness constant, roughly; in fact, it is going to be a bulk modulus.

Here I have taken 2 derivatives of free energy and I have gotten down to bulk modulus, because that is what the free energy is, but when I look at the actual form of the free energy and take 2 derivatives, you see I get a piece like this.

I have one thing that is $H_c d^2/H_c dP^2$. The other piece is dH_c/dP quantity squared -- this is just simple calculus. This is a positive definite quantity. This at T_c is 0. Got it? H_c at T_c is 0, therefore this term is 0. But this is something squared, so it is, in fact, non-0.

I have the utterly surprising fact that at a superconducting phase transition in 0 magnetic field where absolutely nothing happens the atomic positions do not change, you cannot see hardly anything change in the electronic distribution, but the system is now in a different free-energy surface.

At that phase transition I have a step discontinuity in an elastic stiffness, so that by measuring elastic properties at even a weak second-order phase transition, I have a discontinuous physical quantity and, therefore, something you can really sink your teeth into in a measurement.

The same goes for specific heat. There is, in fact, a specific heat discontinuity and I might mention that the first critical test of the Bardeen-Cooper-Schrieffer Nobel Prize-winning theory of superconductivity was an ultrasound measurement. It was, in fact, the ultrasonic attenuation. The freezing out of electronic carriers produced a very strange, exponentially-decreasing-with-temperature form to the ultrasonic attenuation.

DR. QUATE: Are you talking about attenuation or are you talking about bulk modulus?

DR. MIGLIORI: I was talking about bulk modulus here, but that was a slight non sequitur in the sense that I was selling acoustics to this room rather than specifically discussing bulk moduli for that measurement.

Before I get into the measurement techniques themselves, I am going to quickly go through some of these. There is a book on this subject, called Resonant Ultrasound Spectroscopy, and there will be a bunch of pages xeroxed from it for you here only because I am too lazy to extract the figures.

[Transparency 6]

We are going to be talking about resonances and we have seen simple mass-spring resonators described before. The one I am going to deal with -- actually, I am not going to deal with it too much. In terms of the physics measurement, as we get to it later, I am going to want a free resonator, meaning it is not attached to the lab frame.

That is why you will see pictures with 2 masses. This is something that I can let go of in vacuum in 0 G and it will oscillate and be damped by the damping constant associated with this dashed plot and the restoring forces of spring K.

I am going to skip through some of these bits and pieces here until we get to roughly the eigensolutions.

[Transparency 7]

Once you set up the equations of motion for that object it turns out there are magic frequencies. The magic frequencies solve the equations. We are looking at solutions of an equation like this, which really came from a differential equation like that, and we have frequencies, natural modes of vibration of the damped simple harmonic oscillator.

You notice there is a complex quantity. There is a piece here and there is another piece here. I will pick at those in a second, but the point is that if I have a simple resonator with damping, then the frequency is not quite determined by the spring constant and the mass, and we will discuss some vagaries of that in a second, and the damping constant intrudes a bit on one's understanding of the stiffness and the density, or the mass, of the resonator.

The smaller the damping, the less it intrudes, and the better resonant ultrasound spectroscopy, as we get to it, will work but, more importantly, if the damping -- if we are thinking of now a solid instead of a mass and spring -- becomes large, the elastic stiffness is not a good definition or measure of the response of the solid to stress and strain.

As the damping becomes very large, it is disastrously wrong to try to treat the solid as if it were elastic; you might as well just go ahead and treat the solid by its response or solutions to the equation and forget the concept of elasticity. Another place we are going to be is in the region where, roughly, the damping is small (I will define that later).

[Transparency 8]

A measure of the damping is the Q and the Q is basically the frequency of a resonance divided by the full width at half-power points on the resonator, but it is also ω_0 times the K times constant/2 (there are other ways of defining it). This number we will want to have fairly large for most of our measurements.

At a Q of about 100, which means a fairly narrow resonance, we begin to be happy, and when we start to do really precise work on superconductors or silicon or whatever, we will hit our instrumental Q's of 30,000 to 50,000 or so.

[Transparency 9]

What does the resonance look like? What I mean by that is how does the amplitude of displacement of that resonator vary as I drive it with different frequencies. This is a resonator with a Q of 10 and a resonance frequency of $10/2 \pi$, and here is 2π frequency -- there.

I can look at the absolute value of the displacement and I can look at the real part, that is, the in-phase, and then the quadrature piece. The square root of the sum of the squares of this

piece plus this piece gives me the solid line, and that is the general shape of the mechanical resonances that we are going to deal with.

[Transparency 10]

Once again, harping on the concept of Q , if the Q is very large, I have no problem understanding what the frequency of the resonator is, but as the Q becomes smaller and smaller, I find that there are 3 interesting quantities that are left over.

One quantity is the frequency at which the displacement is a maximum and that turns out to be ω_0 , which is the square root of a stiffness divided by a mass, or the square root of an elastic constant divided by a density, times $1-2Q^2$ square root. That is the frequency at which I get maximum displacement.

Then there is the solution of the equations of motion in which I have the real part of the eigenvalue, which is another way of defining the frequency and, finally, there is the frequency at which the motion of the resonator is exactly out of phase with the driving force. You see these are all different and all perfectly good definitions for resonance frequency. As the Q degrades, I have trouble determining what the resonance is.

DR. QUATE: What is held constant in those definitions?

DR. MIGLIORI: Everything. All I am doing is telling you that I have a real part of the eigenvalue, I have the frequency at which maximum displacement occurs, and I have the frequency at which the motion is out of phase with the driving force. Those are all slightly different if dissipation is present and if I am going to attempt to define a frequency, I start to have trouble as dissipation goes up.

MR. WAKELAND: There is a driving force and so the driving force is now $F(\omega)$ is constant? The amplitude of the force is independent of frequency?

DR. MIGLIORI: Yes, of course. I am going to gloss over a lot of these minor points. Some of these you should just assume that it probably was done approximately right.

[Transparency 11]

With that introduction to resonances, let's talk a little bit about resonant ultrasound spectroscopy. You will start to see viewgraphs here that I have used elsewhere, especially as it goes from monochromatic with a simple type face to full color with complicated type faces you can see it is going from colleagues up to funding agencies in terms of the intended audience.

I am going to do this talk in a spiral and start with motivation, I will start to tell you a little about resonant ultrasound, and then I will tell you in more and more detail about especially things that are interesting, so I can stop the talk at any point and you will walk away -- you will walk away, right. (Laughter)

Resonant ultrasound probably -- although Chris Fortunko and Logan and I are really working to make sure we believe this -- provides the highest absolute accuracy of any routine elastic modulus measurement technique. It determines the full anisotropic elastic tensor (that mess I showed you) in a single measurement, so that it makes it very convenient to study objects with lower symmetry than isotropic.

It can handle very small samples -- I am leaving out the optical techniques. I do not know if any of you have ever tried Brouillon scattering, but at 10 K it becomes very difficult. There are other techniques that can operate on small samples, but few of them are as versatile as an ultrasound-based technique, so I am leaving out the weirdos.

It can handle very small samples and, therefore, I have a chance of getting very good single crystals from weird materials. It is very much easier to grow a 1/2-mm single crystal than it is to grow a cm³-single crystal with no defects in it.

Finally, Los Alamos, for lack of any good reason, has probably done more resonant ultrasound than other places, so we get to talk about it more.

[Transparency 12]

One of the really good reasons to use simple mechanical resonances to measure elastic moduli has to do with this. This is a big neutron-scattering facility's interpretation of an ultrasound measurement. What they have done -- this was done at Brookhaven on a single crystal of La₂NiO₄, which is a high-temperature superconductor with a reasonable transition temperature of -- what the heck was it? I cannot even remember it any more -- around 40 K.

What they are measuring is the dispersion curves for the 6 independent sound waves in this tetragonal crystal. They do this by inelastically scattering neutrons off of molecular vibrations, if you will, so they can get an initial slope, which gives you the sound velocity, in various crystallographic directions.

Of course, they do this by measuring the neutron scattering, the number of neutrons scattered into a physical volume in space, which then corresponds to a volume in momentum space in the crystal, so they map out the entire dispersion curve.

The problem is that as they start to go along a body diagonal of the crystal structure, they somehow have to guess that some of these points correspond to this dispersion curve, and some correspond to that, and some correspond to this, so pretty quickly they are in deep trouble unless they already know something, and the precision is exceedingly low.

If I want to see, for example, a 5-ppm shift in elastic module in tin as it goes from normal to superconducting, I have no chance of doing it with neutron scattering, but this technique provides you with data that have the full elastic symmetry in one measurement, so in some sense it is almost the only competing measurement, even though, using conventional time-of-flight measurements for sound velocities, you can get that information pretty well (it takes you a long time).

[Transparency 13]

Another quality of resonances is that if there are elastic-scattering centers present, that is, things that scatter sound without dissipating energy, then in a resonance measurement it does not affect the Q, i.e., let's say I do a pulse-echo measurement and a resonance measurement on some solid object that has a lot of defects in it that scatter sound but that do not dissipate energy.

When I try to run the pulse through it, the pulse hits these scattering centers and diffuses away, so that the received signal is attenuated. That is not thermodynamic attenuation; it is scattering. If I wish to know thermodynamic attenuation, I have to somehow correct for the diffuse scattering.

In a resonance system none of that acoustic energy escapes from the resonator and so it does not see the scattering as a dissipation term, although it will see the scattering as a change in the resonant frequencies. Resonance ultrasound also has a very nice quality of ignoring things that do not have to do with thermodynamics.

[Transparency 14]

Finally, this is the sales pitch for doing resonances. This is a reasonable take at comparing signal-to-noise problems associated with a pulsed measurement and a continuous measurement. The impulse, for that we mean generating a pulse of ultrasonic energy, sending it through a crystal, having it bounce back, and measuring the time of flight.

In the pulsed system I have a huge amount of drive- power-per-unit bandwidth, so that the peak power is quite high. However, in order to measure the time of flight, let's say the time of flight is 2 usec and I want it to a part in 1000, I need a 2-nsec time precision for my electronics, so my noise bandwidth related to that time precision is more like 2 nsec, so in fact, the peak power over the full bandwidth of the measurement turns out to be relatively small.

In a swept sine measurement the peak power-per-unit- bandwidth is not infinite, because I am sweeping the frequency and, therefore, the time I spend at each bit of bandwidth is not infinite and I have to account for the rate at which you sweep the frequency. Typically, for a MHz resonance measurement done in a matter of a few seconds, our power-per-unit bandwidth is probably 10 times higher than the pulsed system.

I already mentioned the noise bandwidth for the receiver. The pulsed system has to have very high time resolution; therefore, the bandwidth of the receiver or its response time has to be very high, so my bandwidth might be a gigahertz, but for the swept sine system the total bandwidth really relates to only the places I need to make the measurement.

If I know roughly where to make the measurement, which is no more of a cheat in resonant ultrasound than it is in pulse-echo and I need to measure, say, 50 resonances to get the elastic constants, my total noise bandwidth can be very, very small.

The duty cycle, that is, the time that I am actually exciting the system, perturbing it so that I can make a measurement, is quite small in a pulsed system, and its continuous, or unity. The detect duty cycle is the same in both.

In a pulsed measurement, the time I spend knowing that a pulse is not there is just as important as the time I spend knowing the pulse is there, so I get a 1 for both of those, and if I take the square root of all the factors, you see that my signal-to-noise ratio in principle is very much better for resonance. Also, it is cheaper to design the electronics and we could go on and on.

[Transparency 15]

Before I describe the details and the catch -- there is a big catch -- let me tell you how well you can do with resonance measurements. I will hit you with these every once in a while and there will be a quiz at the end in which we go over all the numbers, but these will get worse and worse.

This is a measurement of the elastic moduli of single-crystal copper -- ignore the polycrystal measurements for the moment -- in which we have taken our measurement and a literature average here. You can see there are 1.2% differences here; however, in general, there is some reasonable agreement. We can do better than that; I will show you later.

[Transparency 16]

Because the resonance measurement acquires a lot of information about the solid, including symmetry information, we can do the following surprising thing. Here is a single crystal of tantalum and here are the elastic moduli measured by several different groups -- ours are at the end here.

You can see that our numbers correspond very well to the very best values of tantalum (I just tell you that those are the most accepted ones). However, there is one other interesting thing in the fine print. In this case, we did not know the crystallographic orientation of our single crystal of tantalum, so we allowed in our analysis of the data another degree of freedom, and that was the orientation of this cubic single crystal, which we later checked with x-ray diffraction.

We were able to find not only the elastic constants to exceedingly good accuracy but also the crystallographic angles of orientation to about a degree, which is quite good. In order to do that, you have to understand a mechanics problem, which we will get to.

[Transparency 17]

Also as a consequence of the ability to make simultaneous measurements, we can see interesting things in structural phase transitions. For example, once again this is the superconducting single crystal of LaSrCuO. Its crystal structure is tetragonal, which means it is a unit cell that has a square base and an elongated vertical direction and everything is parallel and perpendicular.

When it hits 223 K, it undergoes a structural phase transition, which later, after the break, I will go over in some detail, because it will tell you how we extract very interesting microscopic details from these measurements.

It undergoes a structural phase transition in which one of the 2 shear moduli -- this has 6 elastic moduli, by the way -- collapses, the other does not change. Analysis of our resonance data show that, indeed, 1 shear modulus unchanged and the other collapsing through this phase transition.

It gets very soft. What is happening, very briefly, is that one of these shear moduli is associated with these oxygen octahedra and, as you cool, the octahedra have too much room, because their vibrational energy has decreased. Finally, they just sag -- clunk. When they do that, the symmetry changes and there is a shear mode that causes this distortion if I were to apply an external force.

When they clunk, it means I got a shear strain with no stress and, therefore, the modulus briefly went to 0. It should go to 0 and, of course, because there are very tiny temperature gradients in the sample and very tiny gradients in strontium concentration, it does not go to 0 every place at once in the solid. We lose our ability to fit the data at about this point. This is a measure of the imprecision in the construction of the sample rather than the technique.

All together, resonant ultrasound has become reasonably interesting and important, so that at this point we have people building commercial systems. This evening I will set up a system back there with a sample and it will be there for people to play with and measure.

We have Windows 95 software, in which the entire analysis and taking of data can be run from a mouse. Nobody can read keyboards any more, you have to do it with a mouse. There is a textbook [see Transparency 18] on the subject by some obscure authors and there is a fair amount of literature and lore on how to do the measurements and make the apparatus that you need for low-temperature work.

[Transparency 19]

A lot of this really started with Orson Anderson, who was at Bell Labs and then finally wound up at UCLA. When he got the first lunar finds, which were little particles of sand from

the lunar surface, there was the question as to whether an aluminum silicate found on the moon had the same properties as aluminum silicate or something else on the earth.

Many of these samples that he got turned out to be sub-millimeter spherical samples. It would be very difficult, using any conventional ultrasound technique, to make a measurement of their elastic properties. Orson and postdocs, including Harold Demarest, developed techniques for analyzing the easily measured resonances to extract elastic properties (I will discuss that at some length).

This resulted in this *Science* publication in which he did indeed discover that lunar rocks had anomalously stunningly low elastic moduli. For example, one of the samples here had a sound velocity of about 1.4 km/sec, which matched very well Swiss cheese, shown here at 2.12. Then we have provolone and Romano matching pretty well some of the other lunar rock samples. (Laughter)

These are mineral samples, solid glassy spheres. The answer to the puzzle, of course, is that these rocks were vacuum-cracked on the surface of the moon. Meteor impacts, whatever, produced cracks in these small samples that did not heal because of the vacuum and, therefore, they were very much like a sponge. As soon as they were exposed to water on the earth, the cracks healed and the moduli returned to normal values.

[Transparency 20]

That prompted Orson to send out this Christmas card. "SEASONS GREETINGS/It brightens the spirits/in times like these,/to know the Moon/is made of cheese." Orson is a famous World War II fighter pilot and cartoonist as well as geologist. Included is the extract from the *Science* article when your children ask you what the moon is made of.

MR. WAKELAND: So it was basically bits of sand that NASA picked up and put in some little container that did not let any earth air in and took it home, is that what you are saying?

DR. MIGLIORI: Yes.

[Transparency 21]

You cannot give an ultrasound talk without quoting from Rayleigh, right? So here is the Rayleigh quote.

The catch is taking the resonance data and analyzing them to determine the elastic moduli. When I squeeze a solid like this, it bulges. As the aspect ratio gets close to 1, the kinetic energy and potential energy in the bulging becomes comparable to or greater than the energy associated with the squeezing -- these are all highly technical terms, of course.

You cannot -- you cannot -- simply solve in a 1D way for the eigenfrequencies in terms of the stiffness. It becomes a full 3D problem. To quote Rayleigh here: "In the case of a short rod and of a particle situated near the cylindrical boundary, this lateral motion would be comparable in magnitude with the longitudinal motion, and could not be overlooked without risk of considerable error."

And finally Rayleigh says: "The problem of a rectangular plate whose edges are free, is one of great difficulty, and has for the most part resisted attack." However, he did not have a really nice computer. As I give this talk, I get to cross out one and add another here.

When we started to do these computations -- Los Alamos seems to believe that it can design thermonuclear weapons with computer codes. It is, of course, not true, but they spend an enormous amount of money on big, fast computers. When we first did this, we set the problem

up on a Cray XMP and it would run about 5 seconds per iteration to do this complicated computation, which we will describe, to extract elastic moduli.

Then we tried to do it on an IBM P.C. A.T., an 8 MHz P.C. A.T. and it took 12 hours per iteration, so we threw up our hands in disgust. More recently -- actually, not that recently -- I put this code on this very fine 233-MHz desktop computer I have with a nice 32-bit Fortran compiler and it runs 10 seconds per iteration.

You can see that the improvement of computers is an enabling technology for some experiments. This is one of them.

MR. WAKELAND: In the context of that last slide, what is THE problem of a rectangular plate?

DR. MIGLIORI: I will tell you in a little bit.

[Transparency 22]

The problem that we have to deal with here is to compute the eigenfrequencies of a rectangular parallelepiped from the full elastic tensor. When we do that, we find that it is quite a nasty bit of computation. These are exaggerated distortions that occur for the normal modes of an isotropic solid and each of these is the lowest harmonic of a particular symmetry class for the oscillations.

For example, if the object is a rectangular parallelepiped, then it has mirror symmetry across this plane, this plane, and that plane and, therefore, I should have solutions that are even or odd across each of those 3 planes. There are 8 symmetry classes of solutions and I have shown them to you here.

Our task and, in fact, the entire task of resonant ultrasound spectroscopy is to, one, measure the resonances well enough that you have not messed up the measurement -- that is pretty easy to do, sort of -- and then do a computation that is sufficiently accurate that it does justice to an experimental technique that will be reduced finally to measuring the most accurate of all physical quantities, namely, frequency.

I need to be able to compute the answer to this mechanics problem to probably 6 digits or more. That has been the real basis for resonant ultrasound.

I would myself like to take a break for about 10 minutes to let my voice recover, so let's stop.

DR. MIGLIORI: The catch is to compute the eigenfrequencies, which are what we will measure, from the elastic moduli of a solid.

[Transparency 23]

If you were an engineer, which is different from a physicist, because only engineers have won 2 Nobel Prizes in physics -- John Bardeen, of course, is an electrical engineer -- if you are an engineer, what you would do immediately is to utilize a finite element code to compute the normal modes of vibration of the solid; that is, you would break the solid up into little chunks, get all the little boundary conditions right, and allow a big chunky computer to compute away until it got the eigenfrequencies correct.

It turns out that in this particular case even a very big computer cannot get 6-digit accuracy on the normal mode frequencies by doing it that way; that is, experiment clearly beats the computer in terms of precision. I will show you in a little bit how we have actually used the experiment to check the computer codes.

I say this only because the first approach is even what we thought about doing, just grab a finite element code and go for it, and that is very nice, because those codes are available. For complicated objects, shapes that are unusual, they are the only way to go.

There is a trick and the trick involves writing down the Lagrangian, which, as you remember, is the kinetic energy minus the potential energy. Once again, just to give you a physical feel for it -- I do not think you need a physical feel for it, but I will give it to you anyway.

Here I have a track, it looks like a roller coaster track, it is going like this, it has different heights and valleys. I take a ball bearing and I shoot it along the track, so it is moving fast at the bottom and slows down, then moves fast again, slows down, moves fast again. The horizontal speed of that ball is roughly measured by taking the maximum kinetic energy that I started with and subtracting the potential energy from climbing up and down the hills at any one particular point.

The Lagrangian really tells me how much motion is left in that ball at any particular instant in time. That is also the basis of a famous quantum mechanical approximation technique called the WKB approximation, in which one wishes to compute the wave function of a complicated object or a complicated system. As long as the hills are not too steep, meaning you do not get a lot of reflection coefficient from them, that Lagrangian can tell you just what the wave function ought to do and it works pretty well; in fact, you can get an exact answer for the harmonic oscillator using the WKB approximation.

The Lagrangian for a solid is a nice physically intuitive thing that tells you a lot about the motion that is present. Of course, the kinetic energy is the density, frequency squared, displacement squared, $\frac{1}{2}MV^2$. The potential energy is the elastic tensor times the strains, and a $\frac{1}{2}$ in it.

[Transparency 24]

Using Hamilton's principle, you allow the strains to vary arbitrarily on a volume V and on the surface of the sample. You will note that implicit in this is that the sample is unsupported; it has free surfaces. I obtain immediately, on doing this variational technique, the free surface boundary conditions and all of the possible wave equations for the solid.

[Transparency 25]

To deal with those, all I have to do is find a complete set of functions. As it turns out computationally, the functions do not have to be orthogonal, they just have to be complete. I find a complete set of functions and, as you can see from the previous viewgraph -- which I forgot to mention -- I have turned this volume computational problem into a surface problem because of that second equation relating to free surface boundary conditions.

Now my computational problem is a surface problem, it goes up as the surface of the solid, not the volume, as long as I can find functions that produce the displacements that match what is required by that variational technique. For example, a nice complete set of functions, simply $x^l y^m z^n$ -- yes?

DR. GARRETT: I lost how you got rid of the volume integral.

DR. MIGLIORI: I am not going to go over it in too much detail, but this is basically as a result of Green's theorem once you do the variational piece. The variation in the Lagrangian has a volume piece and a surface piece.

DR. GARRETT: Yes, but why does the volume piece go away?

DR. MIGLIORI: At a minimum the only thing that is going to be left is the surface piece. Let me go over this in some detail later. It is a consequence of a Green's theorem attack on this thing.

Having chosen a complete set, my task is to find the coefficients of that complete set to match free surface boundary conditions and wave equations. I wind up having to solve this sort of equation with a constant removed from it or, in fact, setting $L=0$ here. That becomes a very straightforward computational problem. I want omegas and all I have to do is diagonalize what amounts to about a 1000×1000 matrix in the computer, and that is not too bad a problem to do any more.

DR. GARRETT: You put the ordinate [phonetic] part of your system dead center?

DR. MIGLIORI: Dead center, yes.

DR. GARRETT: Okay, then I buy it, because you need lots of motion in --

DR. MIGLIORI: Yes. Actually we picked Legendre polynomials first, which matched the cubic box kind of math. It turns out you get the same accuracy and it takes a little longer to code it and a little longer to compute.

[Transparency 26]

Of course, as Steve points out, it depends on the basis you choose. If you get the wrong basis, you get these normal modes here of vibration. (Laughter)

This was a fairly complicated computer code. I hate computer programs and I love pocket calculators and you cannot do this problem on a pocket calculator.

PARTICIPANT: Is this computer code available?

DR. MIGLIORI: Yes, I will give it away to anybody who wants it. We have cylinders and spheres and rectangular parallelepipeds.

We really wanted to check the code and I did go through every single calculation in the computer in one iteration but, unfortunately, to diagonalize in a 1000×1000 array, there are so many iterations and there are problems with floating point precision and how many bits the processor carries, and God knows what, so what we did is we took -- you know, every dollar spent for solid-state physics is a dollar taken away from nuclear weapons -- one of the machines that machines things that are this big and roundish [gestures] and we machined this object, which is good to about 1000th of an inch on the surface and it is about 8 inches across.

This is a potato. It turns out that the computational problem for a sphere, which we also solved using the same code and the same techniques, is identical to this object. It is a sphere, except every semi-major axis is different, a sort of ellipsoid but every half-axis is different.

We did this on a numerically controlled milling machine and a bottomless flat. This is trivially easy, to use our computer program to calculate the normal modes. We made the object of nice annealed aluminum, computed the normal modes, measured the normal modes, and we had the first piece of the code done.

Unfortunately, it was only the first piece.

[Transparency 27]

Here is the degeneracy diagram. Starting from sphere we can distort one semi-major axis or several and get to various different shapes here, potato, sphere, ellipsoid, egg, and hemisphere.

This also shows how some of the resonances, the degeneracy, is broken by changing the shape of the object.

Unfortunately, what we really wanted was to determine elastic moduli from resonances, not resonances from elastic moduli. It must be some property of the universe, but you can take the elastic moduli and compute the resonances, but you cannot compute the elastic moduli given the resonances. You have to do some silly approximation business where you try something and it is wrong and then you try and figure out which way to go in 9-dimensional or 21-dimensional elastic constant space to get it better.

That inverse problem is really where all the time goes in doing the computation. We used that module that goes from elastic constants to resonances, make a computation and then iterate that -- and those are the iterations, Tom, that we were talking about. We do a first guess and generate an error function (I will show you that in a minute).

[Transparency 28]

Before we get there, we did other things to check our code. For example, this is supposed to be a little rectangular parallelepiped resonator in here. The piece we used was a fine-grain steel about 3 mm on a side. This is a compressional mode piezoelectric transducer and on top is another compressional mode piezoelectric transducer.

We are touching the corners through a body diagonal of that cube, dry-point contact, very light, maybe gram units of force. What we can do, because we know these transducers are sensitive to only motion perpendicular to their surfaces, is we can compute using the eigenvectors that the code generated when we would get a zero response if we were to tilt the sample.

[Transparency 29]

Here is kind of how it goes. I have a rectangular parallelepiped. I am holding opposite diagonal corners. The corner motion might be like this. If I were to tilt that vector down into the plane of the transducer, I would see no response at the transducer. We could measure null lines for each normal mode.

[Transparency 30]

This is a result of that measurement. This is more along the lines of making sure we knew what we were doing in terms of computation. You can see, as I tilt that sample in x and y, there is a line for a particular normal mode in which the response is 0 for a particular resonance.

[Transparency 31]

We could, in fact, get a whole map of this response here. Basically, although it is hard to indicate in this picture, we had complete agreement between computation and measurement.

In the end, we could measure, putting all these computational techniques together, resonances (and I will talk about that pretty soon) and fit them to fairly decent accuracy.

This is a measurement on a piece of steel used in the transmission of a Saturn automobile, a heat-treatable steel with a very fine-grain structure, so it is very isotropic. It is not even close to the best measurements that we can do. This is very ordinary.

If you are doing 6-digit experiments, at least some of your viewgraphs have to be completely covered with numbers and they may look mind-numbing, but I love to look at them.

We measured these frequencies in megahertz on this little object that was about 4 mm x 3 x 2. Then we adjusted 2 elastic moduli, Young's modulus and shear modulus, if you like, to get these fit frequencies. These are the percent errors between the 2.

There is an RMS percent error that gives us elastic moduli -- looks pretty good. There are not any catches, that works.

DR. GARRETT: Is the K a weighting function?

DR. MIGLIORI: The K is a weighting function and we use it because every once in a while there are accidental near degeneracies. Then you get a mode-repulsion problem -- I will show you that later. I may want to throw out 2 close resonances. Also, it turns out, for reasons I absolutely do not understand at a very fundamental level, that the first resonance never fits very well.

[Transparency 32]

Here is a measurement on a strontium-titanate single crystal that is cubic and has 3 elastic moduli. It is absolutely no more difficult to measure 3 than 2 elastic moduli. Once again, we are in RMS error range of 0.5%. Here is our fit.

I have added something at the bottom of this that took us a long time to understand and I do not know if I completely understand it anyway. I am fitting here 3 elastic moduli for the cubic material and 2 out of the 3 dimensions, meaning I am going to constrain the density but I am going to let 2 dimensions squeeze and the third one move constrained so that the volume does not change for a given mass, so in case I have a measurement error on the lengths, I might be able to get a better fit.

Just think about the 3 elastic moduli for the moment. I have an error function, which is the difference between my measured frequencies and my fitted frequencies. I am fitting to 3 elastic moduli, so that error function has 3 dimensions and a height. It is a 3-dimensional surface plotted in a 4-dimensional space.

It has curvature. When I hit the minimum, I have curvatures. I have 3 different radii of curvature at that minimum. In order to determine the errors, I really need to know if I screw up my measurement how far along on the ball can I go before I start to notice that things are going to pot.

We compute the curvatures and we compute the distance you have to move along the principal axes of curvature, which are not, by the way, the elastic constants, they are random directions. I compute the curvatures and we have, purely by experience, decided that when our chi square for this fit goes up by about 2%, we notice right away that things are going wrong.

I plot the 3 steepest curved pieces here. Look at this line here. This whole part has to do with dimensions; let's forget about that for the moment. This first one says that in some direction I can change c_{11} , which is a compressional elastic constant, by 0.13% and get a 2% shift in chi squared, or c_{12} by 0.48, or c_{44} , which is a shear modulus, by 0.1.

If you take the biggest number in each column, that is the error bar for the measurement. Even though I have a 0.6% RMS fit between frequencies computed and frequencies measured, I am really more like 0.13, 0.5, 0.2, or 0.4% accurate on the moduli. We found that to make some sense, more from experience than from anything else. Yes?

MS. ZOU: What kind of transducers were you using?

DR. MIGLIORI: I will show you. We are going to get to more of those details.

[Inaudible question]

DR. MIGLIORI: No, dissipation-free and, therefore, we want a fairly high Q , although a visiting scientist, Pitoshi Ota, and I and others have worked on the complex elastic constant problem and it is well understood and can be handled. That is, you can measure the attenuation accurately in different directions.

Also, on that same piece of strontium titanate, because we measure all the elastic constants at once, we can measure all the elastic constants at once as a function of temperature and watch it go right through the anti-ferroelectrical transition at 108 K. You can see the elastic constants doing their thing, precursors to the transition, you can compute critical exponents from this, and all sorts of nice stuff.

[Transparency 33]

The most precise objects that humans make are ball bearings -- that is, in production, commercially available. The geometry is quite good. If you take something like silicon nitride, which has been reasonably carefully made, it is very isotropic. If I take a silicon nitride ball bearing with geometry errors of a part in 10^5 , I can measure the elastic constants of that by measuring the resonances.

In this case, the errors get pretty wild. This is a 0.1% measurement of elastic constants, i.e., if I take this thing and I hold my hand 3 inches away, infrared radiation warms that ball bearing up and shifts the elastic constants noticeably for this sort of fit.

This is getting to be very good for us, 0.1%.

[Transparency 34]

The way the measurement is made -- once, again, I do not know enough to get complicated, so I tried to do it as simply as possible -- is we have a frequency synthesizer that drives a piezoelectric transducer. The piezoelectric transducer makes dry-point contact to the sample. I have no coupling fluids, I do not need to glue it on.

The reason that I am making point contact is that I want to excite all of the normal modes; therefore, I want to drive the sample at a low symmetry point, not a high symmetry point, to catch them all. Sometimes we do not catch them all.

The light forces preserve free surface boundary conditions as does the lack of a coupling fluid or glue. I touch the other corner, opposite corner, with another small transducer, amplifiers, phase-sensitive detector, computer -- it is a swept sine measurement -- and in present form we get both in-phase and quadrature data from this.

Tom Gabrielson is my hero here. Think about trying to make resonant ultrasound measurements of elastic moduli of lots of materials with samples of different sizes and different transducers. I may be working in frequency ranges from as low as a kilohertz to 8 MHz in a particularly small sample. I need to preserve as much as signal to noise as I can because of the dry-point contact but, unfortunately, I have no idea what transducers I am going to hook up to, what their impedance is, what their dissipation is.

The electronics design problem is, in some sense, a nightmare, because you do not know exactly what it is you are trying to do, because the experimental situation varies. To make it even more complicated, I have long cables. If I am running a 4-K liquid helium experiment I have to have 6 or 8 feet of cable between a 5-pF transducer and the pre-amp.

In the end we have taken the bite-your-fingernails-desperation approach and done the best we could with the pre-amplifier having nothing attached to it and then hope that any of the transducers would not make things a lot worse.

I will tell you that once we did this carefully and paid lots of attention once we got that first stage done -- in fact, there are 4 separate first stages of the pre-amp in the second-generation resonant ultrasound systems, so you can mess things up in 4 different ways instead of just one -- from there on we knew what to do.

The amplifier string is very carefully designed so that the noise after gain is determined by the first stage of the pre-amp and then the digitization is very carefully designed so that even though we have an adjustable receiver bandwidth -- for example, when we set the receiver bandwidth to 100 Hz, the settling time is exactly 10 msec. That is done with digital signal processing and it is very cute and I will show you that tonight. Finally, even though we use 16-bit analog-to-digital converters, we get 21 to 23 or 24 bits of resolution and 117 dB of dynamic range in this system.

Again, it is simple trickery and I will show it to you later.

[Transparency 35]

The transducers -- we are going to be dealing with millimeter-size samples -- we do not want them to be introducing resonances. By the way, the most important thing is to have no signal and only white noise when there is no sample in place. One way of getting signals when there is no sample in place is to have the transducer resonate.

It gets very hard to make transducers that are very much smaller than a 1/2-mm sample. What we did was take a slightly different approach to it. We used single-crystal diamond that has a higher sound velocity than anything. Therefore, the resonances of the single-crystal diamond cylinder of this size were typically above 4-5 MHz and by constructing this sort of structure here, which is mostly diamond, we could keep transducer resonances quite high.

Of course, their Q's were usually lower than the sample, so even as we started to come up on the tails of these resonances we could select which was sample and which was transducer resonance.

DR. GARRETT: Doesn't your choice of transducer substrate have an effect on the price of this device? (Laughter)

DR. MIGLIORI: Is this the Los Alamos price or the public price? The diamonds are 75 bucks apiece.

DR. GARRETT: Is that true?

DR. MIGLIORI: Yes, they use them for windows and all sorts of experimental stuff. In fact, they were a VWR catalogue item. They are really tiny and especially when we said we did not care if they were a little yellow, then they got really --

[Inaudible question]

DR. MIGLIORI: The thing is, if you are going to do this only at room temperature, you might as well use PZT-5A, because its Q is so low that you can always pick out your sample resonances among transducer backgrounds. But if you are going to 4 K, things get really strangely "ringy," even though some things get very dissipative, like the glass. Kapton plastic is basically an orientational glass; its attenuation goes shooting up at low temperatures.

This structure works pretty nicely cold and we tend to use it.

[Transparency 36 – unavailable at time of printing]

Here is a picture showing you an ordinary every-day object for scale here. Here is the diamond cylinder. Here is a lithium-niobate single-crystal transducer. Once again, we chose those because the attenuation is pretty clean and low in them. This is a sample of La_2CuO_4 superconductor before it has been polished. This is typical of what we measure.

[Transparency 37]

We have also done diffusion-bonded transducers for higher temperature work. Here is a silver diffusion-bonded assembly in which we have an alumina backing plate, a PZT-5A element, and an alumina wear plate -- this is more for nondestructive testing. As you can see, that transducer has no damping elements in it, except for the PZT. Therefore, the only intrinsic dissipation mechanism is the loss tangent of the PZT-5A and it gets a little quieter than average. If you do not mess it up with the pre-amp, you do pretty well.

PARTICIPANT: What is the Curie temperature?

DR. MIGLIORI: Curie temperature is around 350° C. Lithium niobate has Curie temperature of 1100 C, but around 800 things start to go to hell because you excite electron holes across the band gap and it becomes electrically conducting and pretty rapidly its RC time constant goes to hell and you cannot see anything.

[Transparency 38]

Orson Anderson has used buffer rods and made measurements at 1800 C, and at 700 C we are able to use direct contact with transducers.

[Transparency 39]

This will give you an idea of what happens as you warm things up. This is that same 5120 steel that we liked for checking our apparatus out. At 38°C we have this resonance spectrum. At 378 C we start to lose it, because the transducer is losing its ability to generate a voltage because it is shorted by electron holes that are generated with temperature.

[Transparency 40]

This is the high-temperature cell. This was constructed completely with stainless steel, tantalum, aluminum, no glue, no solder.

MR. HICKEY: Does the bottom transducer move relative to the upper transducer, so that you can measure --

DR. MIGLIORI: Actually, the top transducer is loose. It is on a springy little tantalum foil lever. I think I have a better picture of a low-temperature cell that is designed just like this.

[Transparency 41]

We are doing measurements that give us 3- or 4-digit results, which means we need to know the sample geometry to 3 or 4 digits or better. We had a lot of trouble with this. Finally, what we did was to abandon all of the fancy commercial metallurgical polishing systems, take a Pyrex glass plate, some ground-steel shim stock from Starret, and some heat-sensitive adhesives that you could melt on a hot plate.

We trapped a sample in this little area here by pushing the shims in, let it set up, polish it flat. Then we could flip it over, get a second surface that turned out to be parallel to the first one to easily within a micron out of 2 or 3 mm.

Now I have two parallel surfaces. The problem was getting them perpendicular. You hold it up like this and you push the shims in hard this way so that they are holding the sides. Then

you flip it over and you have got it. We were able to get samples that had accurate right angles and parts/1000-or-better geometry.

DR. KEOLIAN: Is the shim that perpendicular?

DR. MIGLIORI: Think about a Bridgeport milling machine. It will do about 1000th of an inch over 20 inches. We are looking for 1000th of an inch, which is a part in 20,000. If we get a micron in 3 mm, the Bridgeport milling machine can do that like falling off a log, so it is not that hard to do.

The trouble is, most of the commercial and metallurgical cutting and polishing jigs are not even close, because they are not trying to make 2 parallel faces. Then when you try to turn them perpendicular, it gets worse. It is a trivial problem that is solved with a trivial amount of equipment, but it took us a while to figure out anything that would work as well as the measurement could be made.

[Transparency 42]

Here is the low-temperature cell. This is a parallelogram hinge arrangement here, so that this arm can move up and down, keeping the transducer perpendicular. Jay?

DR. MAYNARD: I thought I would mention that you have to measure the dimensions, also. We found that Starret sold a dial indicator that is good to about half a micron.

DR. MIGLIORI: Yes, that is exactly what we used, a Starret indicator. It is not bad, and then if you float the dimension slightly, you have got it in the fit.

I am going to talk about plutonium in a little bit, too. This is a low-temperature cell. It has a parallelogram hinge. The sample goes in here. This whole thing can go down to 4 K and up to about 350 K.

MR. SMITH: About plutonium, can graduate students do that?

DR. MIGLIORI: Only the ones we do not like. (Laughter)

Otherwise we do it ourselves. No, you cannot get a graduate student anywhere near a plutonium glove box. If you put a tweezers through a glove in that glove box, you will shut down a 400-person facility for at least a week while they run around in little circles figuring out how to make it safer to be there.

The people who are experienced with glove boxes and plutonium do not want anyone inexperienced in these places. In fact, no graduate students are using --

MR. SMITH: Then how does one gain experience?

DR. MIGLIORI: It is a bootstrap process.

You know, it is not even funny, these guys are nuts. It is really serious business. To get into the plutonium facilities you go through a guard station and there are a bunch of guys with machine guns. Then you go through another guard station with a bunch of more-mean-looking guys with machine guns. Then you go through a 2-foot thick bank vault door with a bunch of really mean guys with big machine guns, and then you are into the glove box area.

There is an alarm that indicates a criticality accident. When that alarm goes off (which it never has, thank God), you are allowed to run out of that building without going past all the mean guards, because they have already left by then. (Laughter)

[Transparency 43]

There are a lot of subtleties to the electronics design and there are a lot of compromises because of the real-world problem of making versatile equipment. One of the problems is that

we have to deal with small samples; I cannot get a good single crystal of, say, nickel aluminum that is much bigger than a millimeter or 2, so my transducers have to be small. They are of order of a millimeter and, therefore, they have to have of order of like 5 pF of capacitance.

To do an experiment at 4 K, I do not want my pre-amps down there, because that is difficult to do for a piece of versatile equipment like that. Also, I am going to be running sometimes from 4 K up to 350 K in a single temperature swing. I have to deal with a 5-pF transducer in series or being connected to the outside world with 100 pF of cable or more, so I have the transducer in parallel with the cable capacitance, and I do not like that.

I need to do electronic things about it, and we will talk more about it later.

[Transparency 44]

Basically, the problem I have is this one, and there are ways around it called charge amps. We will look at that a little bit tonight.

[Transparency 45]

I am going to skip a few things here on op-amps.

You are supposed to look at only the equation here.

[Transparency 46]

Finally, I have a system in which I do not need to have a lot of noise bandwidth present, so I need to do the electronics in such a way that I really optimize my noise band for the rate at which I sweep frequency, meaning I want some sort of tracking filter. As I sweep frequency, I want to be looking at only the frequencies that I am trying to measure.

It is very, very difficult to make an analog tracking filter. What you do is you cheat. Tom has already alluded to this. What you do is a really good analog front-end design and then you finally get digital signal processing at the end.

In the process of going through this analog design you need a tracking filter, and the way you do this is you take the frequency you are trying to measure and you beat it with another frequency that is always a constant difference away. It is very easy to generate a frequency F and $F+1$ kHz digitally, so the beat frequency is a kilohertz.

It turns out all my information now is contained in that 1 kHz if I process it. This is called a heterodyne technique. You can get both phase and amplitude information from that beat frequency but, more importantly, what it has really done is it has shifted the entire spectrum by the frequency $F+\Delta F$ in noise space.

What I mean by that is this. Let's say that I have pre-amps with a $1/F$ tail and some other junk floating around and I am trying to measure a 1-kHz signal. If all of my electronics is really dealing with a frequency F and a frequency $F+1$ kHz, where F is order half-a-megahertz, then all the electronics is really processing half-megahertz signals to get out a 1-kHz piece of information, so I have shifted away from all the $1/F$ tails, but I have to have a higher frequency response for the system.

I finally come out at the end with a device called a mixer, which is a mathematical instantaneous multiplier of 2 sinusoidal signals, one of which contains information and one of which has constant amplitude. I get a sum and a difference frequency from the mix.

I can easily throw away the sum frequency. For example, if $\Delta\omega$ is a kilohertz and ω is half-a-megahertz, then the sum frequency is a megahertz and the difference frequency is a kilohertz. It is trivially easy to filter it out. You are left with the 1-kHz piece.

Now, instead of building a tracking filter, I built a filter that operates at 1 kHz and I can do a really nice analog job of that. When I am done with that, I then do a digital filter or detection on that final piece -- I will not talk to you about that now, we will do it later this evening. The point is, it is carefully thought through.

[Transparency 47]

Finally, we get spectra. Here are spectra taken on a little alumina rectangular parallelepiped. Depending on exactly how it is mounted in that transducer, you see some of these modes are coming and going. For example, this one here -- it is very hard to spot it in this spectrum. These were taken by rotating that rectangular parallelepiped on the transducers. I have to be very careful to be sure that I get all the resonances. Even though I have done a good electronics job, geometry may conspire to prevent me from seeing them.

[Transparency 48]

Sometimes other things go wrong. Here is the avoided crossing problem. These are 2 resonances, one associated with a strongly changing shear modulus, the other with a modulus that is independent of temperature.

MR. SMITH: On that prior slide, you missed some of the modes because the corner is on the mode. If you read just that, can you fix it on a different corner and get rid of --

DR. MIGLIORI: If you think about it for a little bit, you will see that the corners are all equivalent. However, if I rotate it on that body diagonal, I will take that null and move it out of the plane of the transducer and pick up the resonance. Or if I tilt it.

This is another situation where, even though you have done perfect electronics and the sample geometry is perfect, you can get a classic avoided crossing. Here are 2 resonances. When they get very near each other, they repel. The repulsion, the frequency, the amount of frequency repulsion is linearly proportional to the coupling in the region of the repulsion, but only quadratically affects the resonances far away.

Far away I am okay, but as soon as I get in this region I will get anomalously shifted resonances that really have nothing to do with much of the physics and I might want to throw them away.

[Transparency 49]

Finally, I want to go briefly through the fitting procedure once again. We are going to measure resonances, we are going to use a complicated code to compute frequencies from elastic constants, then we are going to generate an error function in which we have a weighting factor and a sum of the differences between the measured and computed elastic constants. I am going to try to minimize these sorts of error functions for this whole process by iterating the code.

[Inaudible question]

DR. MIGLIORI: It is a combination of steepest descent and gradient search. This is a nasty, nasty problem. The codes are reasonably adaptive; they change the mix as you get close to the bottom in ways that I have not been able to understand mathematically completely very well, but by watching thousands and thousands of these fits we have gotten ideas on how to change the convergence parameters that work reasonably well. The codes are pretty mature at this point.

DR. LAUTERBORN: In what dimensions do they work? This fitting procedure.

DR. MIGLIORI: The main codes do 9 elastic moduli in the 3 dimensions, so it is 11 dimensions of fitting. That code is easily constrained down to 2. It turns out there are very special mathematical things that happen with 9 or fewer dimensions so that this roughly 1000 x 1000 matrix block diagonalizes and it is easy.

As soon as you go to more 9 elastic constants, you get only 2 blocks and the computation time goes shooting up to about a minute per iteration from 10 seconds. It is a minor point.

[Transparency 50]

We measure some frequencies, we generate a little input file for the codes, and off we go.

What I have in the next few viewgraphs *[see Transparencies 51-58]* here is a fitting session. I put that in here, but I am not going to go over it for you in any detail at all, but the idea there is to see what happens when you miss a resonance and what you need to get yourself out of that trap.

Obviously, if my measurement misses a mode, I am going to shift a whole bunch up for the code. It does not know what is going on and it tries to make a best fit with a piece of physics missing, so you may have to insert a dummy mode, and the code allows you to put in a 0 that says that is a place holder, I did not measure anything there, but go ahead and compute it as if I had.

The next several pages in your books give you examples of how that works, just in case you want to use this technique yourselves.

MR. HARDIMAN: You mentioned earlier different model types for the resonant frequency. In this case you used the amplitude spectrum, right? What happens if you use a different measure?

DR. MIGLIORI: Lately I do not use the amplitude spectrum at all. I use the real part. I do a very crude what is called a first moment fit to the real part of a Lorentzian to find the center frequency and use that as input to the code. We can talk about that later in as much detail as you like, but I will show you that this evening in hardware. Yes?

MS. ZOU: Suppose you have nonmodular. Suppose you are given a pretty good initial guess. How much time does it take to actually fit, to get a final good result?

DR. MIGLIORI: Forty seconds. Yes, a long, long time on the 233-MHz pentium. It is all CPU time.

What I would like to do is take a break now and we will come back in 10 minutes and I am going to go over what you do after you have measured elastic moduli with a nice high-precision technique and what physics you can learn.

DR. MIGLIORI: I want to talk a little bit about phase transitions and electronic states and things like that, and resonant ultrasound.

[Transparency 59]

I first wanted to briefly remind you of what the free energy is looking like at a second-order phase transition, a continuous phase transition.

What I have is a high-temperature symmetric state, symmetric meaning it has a fairly high symmetry, for example, cubic. As I lower the temperature, the free energy develops a flat spot and then 2 minima and it does this in a completely continuous way.

For example, if I take the order parameter to be the ratio of one edge of the cube to the other, the free energy says they are equal and that ratio stays right there until the second-order

phase transition and then one of the edges starts to fall over into this minimum and they slowly diverge with temperature.

A first-order phase transition is one in which I go abruptly from this sort of picture to this one. The order parameter suddenly changes value instead of continuously moving. In the second-order piece what is missing from here is a very shallow minimum that is only a little way from the minima.

It is basically a Ginzburg-Landau picture of how a phase transition is going and it gives you a rough idea of the sorts of physics we are going to look for in a minute.

[Transparencies 60-62]

The next few are detailed calculations of how to look at and use a Ginzburg-Landau formalism to get the changes in elastic constants.

[Transparency 63]

What I want to do is pick on a generic perovskite crystal structure. This is typical of many oxides, for example, high-temperature superconductors, colossal magnetoresistance materials, ferroelectrics, and anti-ferroelectrics, in which we have, in this case, a cubic array of copper atoms (the dark circles) surrounded by an oxygen octahedron in which there is any oxygen atom at each vertex of this octahedron. This is a setup that is ripe for lots of phase transitions.

[Transparency 64]

As I showed you earlier, and we will come back to that, this particular crystal structure undergoes a phase transition in which some of the shear moduli change, or one changes and the other does not.

[Transparency 65]

How does that work? What we are going to do is a little hand-waving geometry. This is supposed to be the top oxygen in one of these octahedra. These are the other oxygens in this octahedra. The octahedra -- here is one, here is another, here is a third, and you see they are connected at the corners, and so on. There are spaces in here where there is nothing.

I can consider that I have a unit cell that might be this dotted line here and that unit cell might be cubic. The oxygen octahedron you might want to consider to be a reasonably rigid structure, which is not wildly accurate, but just take it that way for the moment. At higher temperatures, its thermal motion makes it take up a fair amount of room. The best way for it to sit in the lattice is straight up and down.

As I cool it off, the amplitude of thermal motion decreases and now the bonds connecting the coppers and other structure in this crystal leave it with too much room, meaning it does not know what to do and it is going to try to rattle.

One of them starts to tilt and it tilts as if this whole edge is a hinge. That top oxygen goes this way. When that top oxygen goes that way, it makes this whole row tilt like this. Then this whole row follows it along the hinge so that the whole thing buckles. Do you see that? As it cools down, I get this buckling of the oxygen octahedron.

When they buckle, this distance changes. You can see if this used to be a square and, therefore, this is a square, if this thing buckles so that it is pleated like this, then this edge becomes smaller than this edge and it goes from square to rectangular, and this is a cubic-to-tetragonal phase transition; that is, it went from all the edges being the same to 2 out of 3 of the

edges the same via this buckling that was strictly controlled by a thermally generated tilt of the oxygen octahedron.

MR. SMITH: Row to row, is that a random occurrence where only some of them buckle along the one, I guess, end point?

DR. MIGLIORI: You see, you think of those attachment points as real honest-to-God ball bearings. They absolutely have to form that accordion structure.

MR. SMITH: So it takes with it the following row and the following row?

DR. MIGLIORI: The whole thing goes until you reach an impurity or some other mess. As you can see, this buckling could go either -- the pleats could be running this way or they could be running this way. I can get a twin in which I have a region where some of them buckle this way and the others buckle that way. Then I have a grain boundary or a defect boundary between those regions, but I will get large macroscopic regions in the crystal where all the pleats are running together.

If I have a really good single crystal, I might get the entire thing to go in one direction but, in fact, you would think that that is what would happen if the crystals were absolutely perfect. However, there are long-range forces, or longer range forces, in these atoms beside next-nearest-neighbor bonds.

Fundamentally -- and I will show you this in a little bit -- you can get systems in which that pleating might go only a couple of lattice spacings before it switches phase. That sort of kink or defect is intrinsic to the physics and not merely a defect that you could anneal away thermodynamically.

Let's get back to the pleating. We just got this edge to get shorter than this edge as the octahedron started to tilt. When that happens, you can see that this square becomes diamond-shaped. Is everybody with me? It is very important.

When I take a square and I make it a diamond, it is exactly as if I sheared it. Therefore, the tilting of these octahedra generates a soft shear mode that is Brouillon zone center, meaning it is a long-wavelength soft-shear mode because that pleating extends over macroscopic distances. If I see an elastic-constant shift that shows a particular shear modulus goes soft with a particular temperature dependence, then I can say a lot about a model for this structural phase transition right down to the microscopic atomic-bond level.

However, you need to do a darned good job of measuring.

[Transparency 66]

For example, we did a darned good job of measuring and so did the neutron-scattering boys. This is a fit to 1 over the soft elastic modulus versus temperature -- it is a Curie-Weiss plot -- which gives us a critical exponent for the softening of 1 .

I do not know if you guys have ever tried to fit critical exponents, but it is very easy to get a half or 1 or $3/2$, unless you have at least a decade in amplitude and you have very low noise data. These are experimental ultrasound data points plotted against a critical exponent of 1 .

[Inaudible question]

DR. MIGLIORI: When I have a phase transition and I have a precursor to it, that softening or change will show some curvature and that curvature is often fit by a particular exponent that is determined by the dimensionality of the phase transition, whether it is occurring in only, say, separated planes or is fully 3D determines whether I have a half or a 1 or other numbers.

If I turn on more precise physics, those exponents can get quite odd but this so-called renormalization group theory can make very strong predictions about critical exponents and the observation of one of those critical exponents can tell you a lot about the system.

Unfortunately, if I am looking at $1/T$ -TC that is a kind of Curie-Weiss or critical exponent of -1, or 1 over the square root of that, when you do the fitting, unless you have really quiet data and lots of range, they look the same.

The other part of this story is that you want to pick a critical exponent of a half for this particular system -- I will not go into the details. You get a half if you do a neutron-scattering experiment of the elastic moduli, because there is so much noise. When you do it accurately with resonant ultrasound you get a critical exponent of 1.

There is no simple mean-field Ginzburg-Landau theory that fits that particular critical exponent over that temperature range, a simple classical measurement picked at a microscopic theory of this phase transition and threw it out of the running, leaving you with a little bit of the mystery that has not been solved yet.

[Transparency 67]

DR. MIGLIORI: They can see the elastic modulus soften and they can see it right out to the zone edge but they have 10% error bars and that is the end of that. It is very useful to have those measurements, because you get zone edge effects, and if you get an accurate zone center measurement you can pretty much piece together everything.

Part of the reason we could piece things together was that we had all the elastic moduli simultaneously, so we had the softening ones and we also had stiffening ones to normalize our data to thermal expansion effects and also to show that some of the elastic constants were not changing. That was quite important for the symmetry-derived picture of this thing.

[Transparency 68]

Another interesting point: Remember, elastic constants connect to the free energy, so if I have a metal with a chemical potential μ -- let's see, the chemical potential is the energy, basically, required to add a particle to the system. That is not the same as the Fermi energy and, in fact, you have to be a little bit careful, because people tend to use Fermi energy in semiconductors when, in fact, there are no particles at the Fermi energy -- but the chemical potential μ is really the Fermi energy at 0 temperature and as you warm up in a metal it has a temperature dependence that goes roughly like temperature -- oh, never mind.

That chemical potential tells you, along with temperature, how the Fermi C is occupied and how fuzzy it is at the top and roughly the density of occupied electronic states.

If I have a strange system, for example, a very narrow-gap semiconductor, then I can predict what the electronic free energy should be. Basically I do that by just doing the sum over all the occupied states. It is very simple.

[Transparency 69]

If I calculate the free energy, it is trivial to calculate the elastic constants. We were studying this narrow-gap semiconductor iron silicide that we had beautiful single crystals of that we grew and we believed that it was characterized by very narrow electronic bands around a narrow gap and, therefore, we could take a fairly crude model of this, fill these up thermally, and compute the elastic stiffness of this due to the 2 narrow lumps of electrons about the gap.

This, by the way, is the predicted magnetic susceptibility.

[Transparency 70]

We made elastic constant measurements versus temperature, which should have given us that, and we subtracted out the lattice pieces. What we were able to do here -- our data are shown, they are nice and smooth -- was fit basically a very nice functional form that indicated 2 lumps of electrons and a very narrow band.

[Transparency 71]

From that, which we published, the synchrotron radiation people were then able to go and look for those narrow bands, electron bands, and that gap, and looking more carefully with high resolution they were able to find it, so simple benchtop mechanical resonance experiment was able to find an interesting detail of the electronic energy distribution in a narrow-gap semiconductor.

[Transparency 72]

Finally, I want to pick at a slightly more complicated structural point and that has to do with the following. It relates a little bit to the problem of the La_2CuO_4 problem, where we have pleats that end and then a new direction starts.

In many systems there may be intrinsic microstructure, meaning that that low-symmetry, low-temperature phase is not translationally invariant in an obvious way but that interactions that go beyond the single unit cell force in microstructure, meaning a little bit more complicated arrangement of the lattice.

You can see that in the following sort of picture. If we take lanthanum strontium manganate, which is another perovskite and exhibits colossal magnetoresistance, in the high-temperature phase it is cubic. As it cools down, it becomes rhombohedral. If you look at the crystallographers, you will become instantly confused and go into another field of physics.

If you try to draw this, actually put the atoms down on a piece of paper and put those atomic positions in place and preserve that cubic unit cell, what you realize is that the rhombohedral crystal structure, which I am showing the unit cell here for, is only the tiniest, tiniest distortion of the cubic structure; that is, it is the cube pushed along these arrows squeezed by such a tiny amount that this angle, which for a cubic unit cell -- you see, I can draw a rhombohedral unit cell inside a cubic one -- if this angle is 60° , they are physically identical, but if it is 60.8° , then it is rhombohedral and the crystallographers draw a completely different unit cell, but physically it is really a tiny distortion of the cube.

We realized this and we began to make some measurements on this system, noting that we expected this same sort of second-order phase transition to occur.

[Transparency 73]

In our now approximately cubic system we could tell if c_{44} changed that we would have this sort of picture of the unit cell, and if the diagonal shear modulus changed, which we call c_+ , we would have this picture.

If you look at this one, all I have done is generated a shear here. I think, if I get this right, none of the inter-atomic bonds change length, just angles. Over here I have really stretched it along this direction, so this bond lengthens and that bond shortens.

Depending on which shear modulus goes, I can tell whether I have a change in a bond length or a change in a bond angle. There has been a lot of controversy over exactly what makes these materials work.

[Transparency 74]

We made a resonant ultrasound measurement and as is typical in these measurements we found that one of the low modes, for example, the first one, was almost a pure c_{44} -type shear -- forget what the $_{44}$ means, but it was almost a shear mode. We could also track that another, higher, mode, the fourth one up, was a diagonal shear mode.

[Transparency 75]

We could then look at the variation of elastic moduli with temperature, so here is the diagonal shear and the main shear. What we find out is that through the phase transition we see in high magnetic field nothing happens to c_{44} and in 0 field it does something strange. Basically I am not going to try to tell you what it is because the theorists are still arguing over it.

What we have here is we have separated out microscopically, on the basis of macroscopic mechanical resonances, which bonds were bending and which bonds were stretching in this material, and we can hand this information to the theorists who care about these things and they can fight over it.

[Transparency 76]

Finally, I want to talk a little bit about plutonium and then quit. Plutonium -- you all know its chief application -- turns out to be an exemplary system for studying complicated symmetry-lowering phase transitions, because it has 5 or 6 of them between its melting point and room temperature.

It goes from body-centered cubic to face-centered cubic, orthorhombic to monoclinic and simple monoclinic as you cool it. It was never carefully studied. It radioactively decays. It was extensively studied, but it was not carefully studied. The reason it was not carefully studied is that it decays radioactively and has a 5% isotope mix, so it is very hard to get things that are really pure, defect-free and, therefore, a lot of the physics may be driven by the presence of impurities and defects. It is hard to separate that out, and it did not matter in World War II.

PARTICIPANT: What about World War III?

DR. MIGLIORI: Well, you know what Einstein said about World War IV. "I don't know what weapons we're going to use in World War III, but I know that it will be sticks and stones in World War IV." (Laughter)

[Transparency 77]

I should also mention that every time I do a measurement, or we do a measurement, on elastic properties of a batch of plutonium, they have to take apart the pit, which is the fission trigger of a thermonuclear weapon, remove the sample from it, and that is the end of it as a weapons component, period. I would say that resonant ultrasound spectroscopy is probably singlehandedly responsible for dismantling at least a few thermonuclear weapons. I think this is a good thing.

It would take me maybe the entire conference to go over in detail the properties of plutonium, so I am going to blast through this -- no pun intended. Plutonium has a very high elastic anisotropy ratio. What that means is that if I attempt to shear it in one crystallographic direction and rotate myself 45° and try and shear it, there is a 7-to-1 variation in the elastic stiffness. That is utterly unique among face-centered cubic metals and very weird.

Furthermore, the size of the shear modulus almost equals the compressional modulus, which is very strange. However, these properties were measured with about 3 atomic % of gallium in the plutonium at room temperature to stabilize the fcc crystal structure.

Three atomic percent means that between 2 gallium impurities there are only 2 plutoniums, if you think about it for a while. This gallium that goes in as an interstitial has sucked the plutoniums in and made a mess of the structure, but people are using those data to make predictive models of this complicated metal. I do not like those numbers, I do not really believe in them.

In order to study plutonium, we need to get a simpler picture of everything that is going on and we need to make measurements in pure single crystals with the hope of probably understanding what is called martensitic phase transitions and other related things.

The basic route to understanding this, and I am going to try to get this one point across very briefly and then quit, is the following. Again, you have to ignore the crystallographers. Remember, we said plutonium goes from body-centered cubic to face-centered cubic to orthorhombic to monoclinic to monoclinic. I mean, it is hopelessly complicated if you look at it that way.

[Transparency 78]

But let me look at a face-centered cubic unit cell. Here it is. Do you see that? Forget this piece and this piece on the end. Here is a square. There is an atom in the center of the face. Face-centered cubic is exactly the same as body-centered tetragonal. Here is the body-centered tetragonal. Tetragonal means this is a square. This one is longer and there is an atom in the middle of it. Physically identical.

I am going to pick face-centered cubic as the underlying crystal structure and describe all of the changes in plutonium as distortions of the face-centered cubic structure. Now there is something really weird that happens and this is a generic weirdness that happens in many metals, especially shape-memory alloys and some high-strength martensites.

Let me take that face-centered cubic. Here it is. I am holding it between my hands. Of course, I could have called it body-centered tetragonal with the tetragonal axis aligned this way. Now I start to squeeze it. I am squeezing it and I am squeezing it. This is the stress-strain curve.

I start face-centered cubic. I squeeze it, I squeeze it. When I reach a magic number where this length C and this length A over the square root of 2 are equal, it is now body-centered cubic. I started by calling it body-centered tetragonal or face-centered cubic. I squeezed it in the long direction for body-centered tetragonal to make that long axis equal to the others, so now it is body-centered cubic.

I am sitting here and I am applying stress on a substance that now has cubic symmetry but this face and this face are stress-free. Therefore, it takes 0 force to hold it in that shape. It is called a Bain transformation; that is, when I squeeze this, I kind of go over the top and clunk. In fact, as soon as I get to this point, the elastic constants go negative and it snaps if it is under constant stress conditions.

The first thing I notice is that it would be extremely easy to go from the high-temperature body-centered cubic phase to the face-centered cubic phase by this Bain transformation route, so that is not such a mysterious phase transition.

[Transparency 79]

Here are the numbers. You can see that an elastic measurement through or anywhere near that phase transition is going to tell me exactly when that happens and what is going on.

By the way, here are lattice parameters for the face-centered tetragonal- or body-centered-tetragonal point of view. If I stick with one, say face-centered tetragonal, it went from body-centered tetragonal to face-centered tetragonal. When it goes orthorhombic, all I really have to do is take that face-centered cubic structure and make one edge a little longer and the other a little shorter, a very natural symmetry-lowering distortion that occurs in many materials as they cool off. It is a just a little change in symmetry because there is too much room.

Now I have got orthorhombic. As I cool it down further, things still have too much room and it tilts a little, click. Now I have the monoclinic phase. It tilts again, click, and I have the room-temperature monoclinic phase.

If you stick with a face-centered description, you can go all the way through all the phase transitions without a lot of mystery and at every step of the way ultrasound measurements that we hope to do will tell you whether you have it right or not.

[Transparency 80]

Here is the effect of this gallium that is typically put in plutonium. It sucks in a lot of atoms and it produces shear-like and compressive-like distortions. You can see that if I have something that snaps from one cubic phase to another as I squeeze it and I put in an interstitial that strains it because the interstitial is too small, I am setting this system up for this snap anyway and I may cause very strange metallurgical things to happen by putting that gallium in that would not happen otherwise.

Another interesting problem is that now I have got plutonium and I am warming it up. When I go from the second- to-the-last phase transition, which is face-centered cubic, to the final phase before it melts, which is body-centered cubic, the volume decreases. This is very unusual.

A lot of the plutonium theorists and also theorists who have looked at nickel aluminum and nickel titanium and other high-strength alloys spend time trying to cook up very weird inter-atomic potentials that would allow the volume to contract on heating.

When the volume contracts -- this is a metal -- you squeeze the Fermi surface down and, therefore, the internal energy has definitely gone up, unless there is extremely strange directional bonding or the free energy actually dropped because the free energy is the internal energy minus temperature times entropy minus TS .

It is possible to have this go this way if the entropy term overwhelms everything at high temperatures. Now how could that happen? Body-centered cubic metals have a very soft shear mode commonly, many of them. If on warming plutonium went from a system that had no soft shear modes to one with a very soft shear mode, you could pick up a ton of entropy, because it can populate this low-energy mode and, therefore, the TS term could overwhelm any changes in the internal energy -- this is my own interpretation at this point.

What we need to do now is understand whether there is a soft mode in the phase in pure plutonium when it is fcc. I know I have lost you all, but the point is that alloyed plutonium fcc shows a soft mode at room temperature but at 400 C nobody has ever measured it. We can measure it with resonant ultrasound because we can do little tiny samples and, therefore, we can see the pure material. That soft mode disappears and we pick up a soft mode in the next phase up.

We have almost a complete picture of the metallurgy of one of the most complicated martensites and also some very fine pieces of theory that we can then apply to everything else that we can find, because in almost no other system do all these phases occur in that order.

I think we will get some very nice science from this if resonant ultrasound can be made to work.

[Transparency 81]

There is a polycrystal plutonium sample with a Q of 5700 for one of the modes at nearly half a megahertz. It is going to be like falling off a log. We will get everything.

[Transparency 82]

Finally, I am going to talk about nondestructive testing with one viewgraph. That is the first launch of a Trident II missile: Your taxpayer dollars at work. Look at that, a 570° turn before -- they blew it at that point because it was not in control. (Laughter)

I would have bought all the rest of them for 4th of July.

Quality control has become increasingly important to American industry and aside from measuring elastic moduli of engineering materials to be used in cars or spaceships or planes or boats or whatever. As a form of process control to determine if the raw materials are good, it is also possible to obtain resonant signatures of objects; that is, you simply take a ceramic valve for a Mercedes car and you take a perfectly made one and measure the first 50 resonances.

Then you note that maybe some of the valves come with a chip on the valve seat and that affects these 3 modes. You can very quickly measure those modes in a production environment and look for the signature of a bad part and use it to sort parts without using dipenetrin [phonetic], eddy current, or anything else. In many circumstances measuring a resonance fingerprint can provide you with an alternate nondestructive testing technique that is cost effective.

This is a particular application of it to a silicon nitride ball bearing in which the bearing was improperly made, so it is slightly out of spherical. As we showed before, a spherical solid, isotropic solid, has degenerate modes, but when it is made aspherical, the mode degeneracy is broken, the modes split, and the size of the splitting tells me how much out of round this bearing is. You can use your imagination to think of other ways of doing that sort of measurement.

What I skipped over in there, by the way, are just typical examples of phase transitions and other physics that we have observed, superconducting, anti-ferroelectric, and what. You can just thumb through it yourselves when you need to fall asleep and you are having trouble.

[Transparency 83]

I will end with this. We continually find new things to do with this and interesting things to play with. Does everybody see the bug? See the walking stick? Thanks.

PARTICIPANT: You did mention the small sample size and also the bridge.

DR. MIGLIORI: And everything in-between.

PARTICIPANT: Have you measured things in-between?

DR. MIGLIORI: Yes, high-precision parts that are relatively small so that you can mount them in a free surface way are very amenable to this. Big crummy parts -- for example, an aluminum casting for a C130 airplane landing wheel is about this big, it is a forging, actually, and there is forging-flashing on it that makes every single wheel slightly different, but they are all good, so you cannot test them.

But ball bearings are great, because they are precisely made.

PARTICIPANT: What about a core sample of rock from the earth that may be --

DR. MIGLIORI: There is an extensive effort at Los Alamos to study the resonant ultrasound properties of core samples in other rocks. I can put you in touch with those people.

DR. GARRETT: On the viewgraph that had the 6 phases of plutonium, the number of elastic constants changes as you go through those different symmetries.

DR. MIGLIORI: Yes, but not the number of resonances, because there are no degenerate modes if you cut the sample properly. Is that what you are after?

DR. GARRETT: I am sort of curious to know in each one of those particular cases how many independent modes there were, because you are going to have these resonances, you are going to interpret them as difference c's. There may be 5 c's for one phase and 9 c's for another phase.

DR. MIGLIORI: In fact, there are 3, 3, 9, 13. That is the way it goes. Yes, the number of elastic moduli change drastically but none of the edges of the sample is any magic ratio, there are no degenerate modes and, therefore, all the modes change, the mode positions change, and it is the particular pattern of positions that determines how many elastic constants there are, not anything special. It is the numerical quantitative positions of the resonances that determine the number of elastic moduli. It is a very difficult computational problem because of that, and it is quite tricky. You use every bit of physics you can to help with that.

MR. HICKEY: What determines the shape of the sample that you choose?

DR. MIGLIORI: If I want to use the Lagrangian minimization technique, I need a simple shape where I can do that surface integral. Therefore, we have developed codes for rectangular parallelepipeds, spheres, and cylinders. For anything else we use a finite element code.

MR. HICKEY: Do you use a cube or is that --

DR. MIGLIORI: The cube is bad news, because if you have a low-symmetry object you will have degenerate modes. They will be sitting right on top of each other or they will repel each other weakly and you will have a mess, so you never want to use a cube, which means when you make these samples you do not bother measuring them, just cut them up, make sure the faces are parallel. Whatever they come out with you measure and you go for it.

MR. WAKELAND: So if you have a single crystal, you want to use a rectangular parallelepiped, but if you have a homogeneous material you want to use a sphere?

DR. MIGLIORI: Every single material that this technique is valid for has to be homogeneous. Anisotropic materials need to have the axes aligned with the edges of the rectangular parallelepiped. You always want to use a rectangular parallelepiped if you can, because the computation is faster.

MR. WAKELAND: You showed us a picture of a little piece of plutonium, and what did you say about it?

DR. MIGLIORI: I said its Q was really high and I did not know the dimensions yet. With that high Q it is quite clear I will get very high-accuracy measurements throughout the temperature range on it.

DR. KEOLIAN: Maybe I do not understand it, but why do you need about 50 frequencies measured to 6 digits to get 3 elastic constants measured to 4?

DR. MIGLIORI: Bonzo, it is great having you in the audience. This is a great question. There is a very subtle answer to it. Many of the modes, about 85% of them, depend on one of the shear moduli in the same way. Even though I measure 50 modes, of those 50 maybe 32 of them depend purely on c_{44} , and it is only the other few that depend on the other 2 moduli for a cubic system.

Therefore, I have to keep measuring and checking my fits until I have enough information to fit those other moduli. At high frequencies the reverse happens. Most of them become compressional modes. It is a tricky problem and you have not enough redundancy until you get to a lot of modes. You need, really, only about 20 for 2 elastic constants, 15, maybe. For 9 elastic constants it is nice to have 50.

MR. HARRIS: You mentioned including losses or making the elastic constants complex. That, then, would be associated with some dispersion?

DR. MIGLIORI: No, dissipation and dispersion are not --

MR. HARRIS: If it is a causal system, if there are some losses, there has to be some frequency dependence of the constants. Otherwise, if you put a pulse to it, the pulse would get there before you sent it in. My question was do you have any thoughts about what that does to the --

DR. MIGLIORI: Actually, you lost me. Kramers-Kronig is just fine. That tells me, if I know part of the spectrum, for example, I can use that to give me the relationship between the real and the imaginary parts of the spectrum where there is a known relationship, but it does not tell me that the speed of sound has to be different.

MR. HARRIS: It applies to the refractive index and lighter or faster --

DR. MIGLIORI: Sure, this tells me where I have resonant modes.

MR. HARRIS: If you measure at one frequency, then, and you do have attenuation, assuming this relationship, you will have attenuation that would be accompanied by dispersion, so the question is do you get different elastic moduli, slightly different, at different frequencies?

DR. MIGLIORI: You are worrying about the second-order effect? What am I missing here, Steve?

DR. GARRETT: He is saying the modulus will be different. When you make your, say, 50 measurements, say you are looking at the c_{44} or whatever one that is that gives you the same modulus for a whole bunch of different modes, if there is attenuation, the modulus will be different at those different frequencies --

PARTICIPANT: The modulus itself is frequency-dependent.

DR. GARRETT: You are assuming that the modulus is frequency-independent.

PARTICIPANT: But it is second order.

DR. MIGLIORI: It is a second-order correction and if the modulus is frequency-dependent, ye who enter here abandon all hope. (Laughter)

However, if you make the simple correction that the modulus is complex and you make resonance measurements over, say, a decade in frequency or even a factor of 5 in frequency, then from that and the complex solutions, it is only another matrix that is added into the problem.

You can extract the complex part of the elastic moduli and you can even have a shot at doing it, if it is frequency-dependent, if you kick in a Debye relaxation mechanism for the ultrasonic attenuation. This has been handled, been published. It is somewhat complicated and

you need to be fairly desperate to use this technique to do it when you can pick up a pulse-echo ultrasound system and directly measure the attenuation in many systems without having to do that computation.

MR. HARRIS: That means that frequency also depends on the dimensions of the sample.

DR. MIGLIORI: Absolutely.

MR. HARRIS: So if I had a single crystal in different dimensions, I may actually get different moduli.

DR. MIGLIORI: Exactly. Bob Lesure at Colorado State has used that technique to extract information about dissipation associated with interstitial hydrogen motion in metals.

DR. HARGROVE: Has anyone successfully attacked formulating RUS for visco-elastic materials?

DR. MIGLIORI: No. I have tried to make measurements on things that approached visco-elastic and when the Q gets low, it is just a waste of time. There are too many easy problems to solve. What was it that Ronald Reagan said? Hard work never hurt anybody, but why take the chance? (Laughter)

ACOUSTIC RESONATORS AND THE PROPERTIES OF GASES

Michael Moldover
National Institute of Standards and Technology
Gaithersburg

DR. MOLDOVER: I come here with a lot of humility. We have heard about femtoseconds. We learned about state-of-the-art instrumentation and how to design an experiment from the ground up by Tom Gabrielson and Albert Migliori. We heard about really complex acoustic systems, about bubbles from Professor Lauterborn, and about sound propagation in sand and beds of glass beads. We heard about gigahertz from Cal Quate in acoustic microscopy. We heard about acoustic Mach numbers of 0.2 on up to 1 from Mark, and Steve yesterday told us about acoustic resonators where you put a kilowatt in and these resonators must have some Q, maybe somewhere between 10 and 100, so sloshing back and forth in that resonator is 10 kW of sound, or maybe 100 kW. It is a hell of a lot of energy. Just think of the parameter space. I am talking to you at what? a microwatt?

I am going to talk to you about sound in the kilohertz region. The Mach number, as I have learned to call it, is about 10^{-7} or 10^{-8} . The theory -- Rayleigh would have been happy with all of it, I think I learned it from him. Instrumentation is commercial, we can buy it. We can afford to spend the time and energy to refine it. In terms of the complexity parameter, the acoustic system we are going to talk about is as simple as we can make it.

This program has been going on for a long time. In 1978 or so a measurement of the universal gas constant was completed at the National Physical Laboratory in England, which was sort of a counterpart of NBS, and they got the wrong answer, and the answer was wrong by about 150 ppm. Everybody knew it was wrong but they said we are the National Physical Lab and this is what we got, and it was an acoustic method.

My boss called a bunch of people together -- we were, then, the heat division -- and said, "Can you guys think of another way to measure the gas constant?" The idea of using a spherical resonator occurred to me and I called up Izzie Rudnick. Izzie said, "Yes, it'll work and, by the way, you should go meet Mo Greenspan who is over in the sound building in your lab." Those were the most important 2 pieces of advice I ever got, that it would work and that I should meet Mo Greenspan.

This has been going on for a long time. About 1986 we talked about putting together everything we learned from a 1-m aluminum sphere, and it is in a 20-page paper in JASA. In 1988 we finished the work on the gas constant, and this is in a 60-page paper in the NBS *Journal of Research*. I wrote these papers with my co-authors as carefully and as detailed as I possibly could.

There is no point in trying to condense them into some short set of notes; it would be a waste of time. If you really want to know the details, I think they are there. What I hope to do, instead of conveying the details and the equations and the i's and pi's, is to convey the philosophy with which we do the work, why we do it, and what we are planning to do.

[Transparency]

I am going to start off with some motivation. People think of the Bureau of Standards, now the National Institute of Standards and Technology, as the sort of place you go to, to get you thermometer calibrated, the place you call, as Steve Garrett did, if you want the heat capacity of sulfahexafluoride.

I think acoustics is a very good way to calibrate thermometers and I think it is a very good way to get the heat capacity of sulfahexafluoride and I would like to convince you of that, but this is a competitive world and we have to also look at the competition. Why do you need to send us a thermometer to get calibrated? Why are there temperature scales? In general, I want to provoke the question of why.

[Transparency]

I am going to start off by telling you what the situation is in thermometry today. To me, as a physicist, it is intellectually very messy. The people most concerned about thermometry are physical chemists and they are actually very smart people.

They can take a glass container just about this big and fill it with water, evacuated water, water with the air pumped out, de-gased -- it is filled up to there with water -- and they start this freezing, they cool it to near the freezing point, put some ice in the center and started it to freeze (this is solid water). Then they warm up the center and melt it in the center, so you have liquid water, ice here, and water vapor up on top.

That is a triple point, that on the phase diagram, where you have solid water, liquid water, and vapor water, there is only one temperature where that solid, liquid, and vapor are in equilibrium. At that spot the temperature is, by definition, 273.160 K, and you can put as many 0's as you want. It is about 0.01 C, for historical reasons.

Incidentally, it is not ordinary water. It is SMOW, standard mean ocean water, that has been distilled twice, not 3times, because that will change the isotopic distribution. You can buy this for about \$500.00 and most any physical chemist can manipulate it such that he can get the temperature there to about a part in 10^{-7} (that is 30 uK). The only thing I know that beats it are the watches we wear, which are also accurate to about a part in 10^{-7} , and they cost \$20.00.

If you had to push it, you would start worrying about the size of the ice crystals, the strain, the isotopic composition. You could probably push it another order of magnitude. I do not think we are going to be able to beat that for defining a temperature, not at that price. As I say, this can be done almost anywhere in the world.

Not only can you do this at the triple point of water but you can do it at the triple point of gallium, you can do it at the triple point of mercury, you can do it at the triple point of hydrogen, so there is a whole bunch of fixed points that people can trade around the world and everybody can get them pretty well.

Then you need something to interpolate between these points. The people who worry about this have decided on using platinum wire purified to a certain point, so there is a recipe: Ye shall get a platinum wire of such and so purity and put it in this bath, then that bath, and that bath and then take the resistance ratio and fit it by such and so a formula. That is the way in which temperatures are defined and thermometers are calibrated and it is pretty messy, because no piece is -- yes?

DR. HOFER: For calibrating thermometry you want to have 2 standards, at least for a lot of things, to get the scale right, I would guess. The water triple point is an obvious one. What is the most popular other fit point?

DR. MOLDOVER: All the other points are weighed equally. That is another thing, the scale is not unique. No 2 pieces of wire are the same and if you calibrate the same piece of wire but at different fixed points you get different results, so it is really intellectually messy.

Frankly, where I would like to see us head is that this is the defining temperature and that any other temperature is the ratio of the speed of sound at that other temperature in an ideal gas to the speed of sound of this temperature. That is intellectually clean and that is what I would like to see us achieve.

[Transparency]

There is a particular reason for being interested in thermodynamic temperatures and that is the second law of thermodynamics relates things that you can measure in different ways, like the latent heat of a fluid to the derivative of the vapor pressure and the volume of a gas in a liquid.

There is an electrochemical reaction, there is a heat of reaction, there is a charge transfer in that reaction, there is an EMF in the cell, and the temperature dependence of the EMF. There is the Carnot efficiency, the maximum efficiency you can achieve with any heat engine. All of these relationships, which are built on the second law of thermodynamics, are true only if the temperature in these relationships is the thermodynamic temperature, so an arbitrary temperature scale will not do.

[Transparency]

I have told you we are heading toward acoustic thermometry. Statistical mechanics, or kinetic theory, tells us that the mean squared energy in a dilute fluid is $3/2 kT$ and that can be related to the speed of sound. When we measure the speed of sound in a gas, we are measuring the mean squared velocity of the gas. We can measure this at the square of the speed of sound twice at a temperature that we would want to measure and the temperature of the triple point of water, which is defined to be this number, and that way we will get our unknown temperature.

The way in which we are going to measure the speed of sound is to use resonances. The speed of sound has the dimensions of length over time. The length that is going to be in here is going to be the cube root of the volume of this resonator. The time is going to be the reciprocal of the resonance frequency, and then there is an eigenvalue that you have to know exactly.

The volume changes as a function of temperature, so for this measurement we need volume ratios and frequency ratios. The volume does not change a lot, so we will be satisfied with measuring the change in volume divided by the volume; we are not trying to measure the volume entirely. Of course we can measure frequency ratios as well as you like.

[Transparency]

This acoustic thermometer is only one example of several competing thermometers. These are thermometers that purport to get to the thermodynamic temperature scale. One knows the equation of state of a dilute gas, $pV=RT$ with correction terms. That is one of the basic thermometers used today.

Here is our speed of sound. At high temperatures, where it is not easy to do gas thermometry, one uses blackbody radiation. Noise thermometry is very useful at very low

temperatures and, under certain circumstances, very high temperatures. These are all absolute thermometers.

These absolute thermometers are based on thermodynamic systems that we understand exactly, we think we can calculate to the level needed, their properties, from quantum mechanics and statistical mechanics. It is this exactness that enables us to use them to define temperature or measure the temperature.

You notice that the temperature is always next to a constant that relates temperature to energy, either Boltzmann's constant, if we are talking about individual particles or individual modes of an oscillator, or the gas constant if we are talking about a mol of particles.

There is a constant there that is important to only thermometrists. It comes up if you have, let's say, a thermometer that measures blackbody radiation but you cannot calibrate it at the triple point of water. Radiation is too weak at the triple point of water to calibrate it well enough, so you want to calibrate this thermometer or use it at some high temperatures and relate it to acoustic thermometry that works near the triple point of water, so we have to know this constant to connect different thermometers, the gas constant or Boltzmann's constant.

[Transparency]

The thing that got our management's attention a few years ago was the claim that we could use acoustics to measure the gas constant. That is the same spherical resonator I have already shown you.

[Transparency]

We are going to fill this resonator -- we are going to use the same relationship relating the speed of sound in a gas to the temperature. We are going to do it at the temperature of the triple point of water, but in addition to measuring the speed of sound we are going to solve for R and we are going to measure the molar mass of the gas we know, so we are really going to have to know the gas, not that it is not sufficient that it be monatomic, but we have to know the mass.

We also have to know the speed of sound in meters per second, speed-of-sound ratios that we use in thermometry. Here we need to know the speed of sound in real S.I. units. We have to know the volume in meters and the frequency in inverse seconds.

We are setting up ourselves to measure the resonant frequency of a cavity at the temperature of the triple point of water. We are going to measure the volume at the temperature of the triple point of water. We are going to take these things in the limit that pressure goes to 0 -- after all, both the volume and the frequency will depend on the pressure.

We are going to measure the molar mass of the gas and, of course, we have got to know the temperature. Yes?

MR. POESE: I do not understand. How do you know that the time average of V^2 is $9/5 C^2$ is the speed of sound?

DR. MOLDOVER: Kinetic theory tells me that. There is some number there, okay, it seems reasonable to you that there is some number there?

MR. POESE: Yes. Are you going to fill these spherical resonators with water?

DR. MOLDOVER: No, sir, I am going to get at the temperature of the triple point of water. I am glad you asked. That temperature, because people can reproduce it to 10^{-7} for \$600.00, we are stuck with that. It is a definition.

But we want to measure some other temperature. We are going to measure the speed of sound in a monatomic gas at the other temperature and at that temperature. That will be the way we will calibrate thermometers. I am glad you asked, and now we will get into acoustics. This is the motivation.

[Transparency]

We are going to look at radially symmetric resonances in a sphere. We started off thinking this was important because you have the highest possible Q. The gas is moving only in and out, there is no viscous dancing at the walls. That turned out to be not particularly important.

It turns out that you do not have to make a very good sphere, and I will tell you more about that, in order to do this well, and that turned out to be very important. The radially symmetric modes are nondegenerate, and that turned out to be of overriding importance.

[Transparency]

This is the apparatus that was used to measure the gas constant. It is a resonator, 9 cm in diameter inside, and it has got walls about 3/4 inch thick, steel. It is about the size of a volleyball. It has got a thermometer here and a thermometer here -- and a thermometer here.

It has got a transducer here, a transducer here. One of them transmits, the other receives. It has got a valve for letting gas in, and so forth. We will talk about what we have to when we need it.

[Transparency]

Here is a picture of the beast. We tried high-tech methods of electron-beam welding the 2 halves together. In the end we wound up using a high-temperature wax, L wax, so this really was done with sealing wax.

[Transparency]

One of those microphones transmits and the other receives and you get a whole spectrum of resonances: Here is one, here is one, and so forth. All these resonances are determined by the speed of sound in the gas with small corrections that have to do with viscosity and thermal conductivity.

There is a tremendous amount of redundancy here. The redundancy is of great value, because it helps us figure out what we are doing and then trying to convince you we know what we doing. This is kind of the reverse of RUS as applied to crystals, where there were 9 unknown elastic constants and 3 unknown dimensions and they measured a whole bunch of resonances to home in on those unknowns. Here we know everything and we are using the resonances to possibly identify our points of ignorance.

[Transparency]

If you take one of those resonances and use a lock-in amplifier and sweep through it, you see an in-phase response and a quadrature response and you fit it to a Lorentzian.

[Transparency]

This is one equation you have to know -- and it will be on the exam. You find the deviations are less than a part in 1000 of what you started with, which means that you can split the line shape, find the center of that line to a part in 1000. This was an imperfection in the lock-in amplifier, and nowadays you can do a factor of 10 better, if you have to.

We are talking about lines of a Q of about 5000, so you can measure the center frequency to about a part in 10^7 , which is much better than we can imagine needing, so the question will not be how well can you measure the frequency but how well can you understand it.

This, incidentally, is the definition of the resonance frequency for our purposes, the experimental definition. Albert mentioned that there are several different definitions.

[Transparency]

You could take the place where the phase shift was 0, or the maximum amplitude. This particular definition, where you fit the complex voltage to a Lorentzian with a complex frequency and complex additive terms to allow for the other resonances, this particular definition, if you have a transducer that has a phase shift that is a linear function of frequency, that will be taken out by this fitting process.

If your transducer has a diaphragm resonance some place and you are not too close to it or you have a phase shift in your electronics, this definition is immune to that.

[Transparency]

In the spectrum I identified the modes with 2 numbers just like the quantum numbers that Jay told us about Sunday night. We are solving the same equation with a different boundary condition. His wave functions had to vanish against hard walls. In our functions their derivatives have to vanish against the wall. We have a standing wave in there. The solutions are spherical Bessel functions and they are different spherical Bessel functions, 1, 2, 3. The first quantum number is which spherical Bessel function we are talking about.

This equation has several roots for a given Bessel function. We could have had the wall there and satisfied the boundary condition -- or the wall there. We also order the roots. It is not exactly the same notation that was used in Jay's talk, but it is the same idea. That equation and those solutions you will find in Rayleigh.

[Transparency]

The most important corrections to those frequencies we get by perturbation theory. Jay mentioned perturbation theory. He was thinking of perturbations that occur throughout the volume of the resonator, or the atom. Most of our perturbations are on the boundary. This is an example of boundary perturbation theory, where the integral is over the bounding surface, the wall of the resonator. There is the wave function squared, and there is some admittance, which is the perturbation, and there is a normalization, or the integral of the wave function squared over the volume.

The admittance for heat flowing in and out of the boundary, the sorts of things that Jay talked about yesterday, has equal real and imaginary parts. It depends upon Γ^{-1} whenever we have heat flow, the heat capacity comes in, in this way, and it depends upon the thermal penetration length.

We can integrate the solutions to Bessel's equation, Bessel functions, over the boundary and calculate the perturbation for this admittance, the thermal admittance of the boundary, and this is the answer you get. It depends upon the eigenvalues of the mode, and n is the index that goes with non-radial modes. For the radial modes n is 0.

[Transparency]

I hate integrating Bessel functions, so I have given you a back-of-the-envelope calculation, where we simply ignore the fact that there is a Bessel function there and we say that the energy is

spread throughout the volume uniformly. The Q is 2π times the energy stored and the energy is dissipated within a boundary layer of the surface.

The volume in which there is dissipation is the area of the sphere times the thermal penetration length -- that is 4π times the radius squared times the thermal penetration length, Γ^{-1} because the thermal effect depends upon Γ and when Γ is 1 we have no heat change. That is the dissipation in a cycle. Here is the volume of the sphere and here is the 2π . You get the right answer except for this factor of $2\pi/3$. We will use this sort of calculation later on today.

[Transparency]

Here is the standing pressure wave, in blue, the thermal wave that tracks it, in green, except at the boundary, we blow up the boundary. Heat is flowing in and out as the pressure goes up and down -- you saw that in animation.

One calculates a frequency shift and a contribution to the half width of this amount, and now I tell you the most important thing about perturbation theory that you have to know, and that is that perturbation theory works. It works very well.

[Transparency]

Here is a measurement of the width of the first, second, third, fourth, fifth radial modes at some fairly low pressure, half an atmosphere. This is the calculation that we just went through, the boundary losses from heat conduction. This is the bulk loss in the gas, the thermal conductivity and viscosity losses, which become important at low pressure.

We calculate the total losses and you see that the measurements are very nearly on the curve. They deviate above the curve a little bit, and that is important, because if they were below the curve it would mean we were creating energy. These losses have to be there and then there are other losses that we have not accounted for. Those extra losses are a measure of our ignorance.

[Transparency]

One of the things we like to do is measure the excess half widths, which we call $\delta G/f$. We measure them in parts per million because we are shooting for parts per million in the speed of sound. Here is a plot for the first radial mode, which is labeled 0.2, the second radial mode, labeled 0.3, and so forth, as a function of pressure going up to -- this is probably 5 atm (the scale got lost).

What you see is for these modes the excess half width is less than a part per million of what was expected. In terms of the width itself, the fractional part of the width, we can measure the thermal conductivity this way to something better than 1% at low pressures.

[Transparency]

More perturbation theory. You can apply perturbation theory to the shape of the resonator. We had 2 hemispheres made for us in the shop and it turned out that they had different radii when we got them. They differed in radius by about 1 ml out of 3.3 inches inside, which is about 3 parts in 10^4 .

You can, or Jim Mehl can, calculate the perturbation to the eigenfrequencies from this shape imperfection on the assumption that it is a smooth variation in shape, so we do not take into account the small step here or we smooth it out.

This is the perturbation. What you see is the wave number is changed something like a part per million, even though the error is 3 parts in 10^4 . To first order the wave number for the spherical mode depends upon the volume and only to second order, that is, in the square of the perturbation, does it depend upon imperfections, or deviations from a spherical shape.

That turns out to be the key to this experiment. Izzie knew this. He told me about a paper in nuclear physics. There was a shell model for the nucleus and you take a nucleus and distort it from a sphere to an ellipsoid and the energy levels in the shell model changed. One can show the energy levels in the shell model depend upon the deformation squared, and the linear term cancels out.

Here is another example of shape perturbation, where we distort the sphere into a spheroid, this radius being reduced by the square root of $1+S$ and this radius being increased by $1+S$, so the volume is unchanged. The eigenvalue for that perturbation, also, with $S \approx 3 \times 10^{-4}$, gives you this kind of quadratic dependence on frequency or mode number.

[Transparency]

I mentioned the redundancy in the spectrum and I mentioned that we are going to extrapolate to 0 pressure. These were sort of the first data we got where we tried that. We are looking at pressures up to about 10 atm. We are looking at various radial modes and we are measuring how their frequencies vary with pressure.

If we understood what we were doing, all these points would be on one straight horizontal line, so we have left something out. What we left out varies from 0 to 300 ppm for this mode and from 0 to +100 ppm for that mode, but they do look sort of like straight lines, so let's plot the slopes and the intercepts, because we are interested in extrapolating to 0 pressure.

[Transparency]

If we plot the intercepts, most of them fall nicely around 0 ppm, certainly within + or -10, except for one that off the top here at 256 and one that is up here at 30. We plot the slopes and the slopes are negative for the lower modes, very negative. Here is a positive one and then there is one up on the ceiling, and then there is another one that is positive.

I have drawn a curve through the slopes and this curve is a 0-parameter curve. It comes from a very simple model. We take the shell and consider it as a thin shell and it has a breathing mode. If the shell were in vacuum and you hit it, it would ring in a breathing mode. The frequency of that breathing mode would depend upon the speed of sound in that shell.

It turns out you cannot do much about that breathing mode. If you make the shell twice as thick, you make it twice as massive, and you also make the restoring force, which is the stretch that occurs when you make it breathe twice as large. The frequency does not change if you make the shell thicker. You are stuck with that frequency and it is going to depend upon what the shell is made of.

You can win by making the shell thick, because the frequency perturbation on the gas depends upon the ratio of the impedance of the gas to that of the wall or the shell, so one has ρC^2 in the gas, which is the same for all noble gases, you cannot beat that, at some pressure, compared to ρC^2 of the wall.

We could have made it out of beryllium, which has a very high speed of sound, but my metallurgy colleague said beryllium is cubic and it is not dimensionally stable and it is poisonous and stay away from it, so we did not make it out of beryllium.

The coefficient here depends upon a/t , and that is why we went up to $3/4$ of an inch. It turns out that this saturates. If you make the radius very small compared to the thickness you do not drive this down to 0, you drive it down to some finite value, so you just do not win by making it thicker and thicker after a while.

You are stuck with this resonance denominator that depends upon the speed of sound in the gas to that in the wall, and an eigenvalue. At some point this is going to be 0 and that is going to resonate and you are going to have to live with it. All spheres have this; so do cylinders and cubes and whatnot.

One perturbation that, for reasons I will tell you later, we got interested in all over, again, when we went back to look at it, we made a more complicated model that Albert called -- he talked about frequencies repelling each other -- this is a plot of the frequency of the shell as a function of temperature, and it hardly changes at all, the elastic constant of steel around room temperature hardly changes, so that shell resonance goes along. Here is a gas resonance, one of the radial ones that couples effectively, and what one expects to see, as Albert told us, is the gas resonance bending over and becoming shell-like and the shell resonance bending over and becoming gas-like.

[Transparency]

This is a very simple-minded model where we have a light mass and a light spring, damped, and a heavy mass and a heavy spring, damped, and we are just going to tune, for example, the light spring and bring the light mass through. It is sort of fun to look at the response function, because we do not measure frequencies, we measure response functions and we fit them.

I said to Jay, "Gee, this would be something nice to animate, and he said, 'Yes, I can do that. I can do it easy.' How long? 'Two hours.' No, no way." Well, he did it in two hours. This guy is good. [The demonstration was shown.]

There is the resonance in the gas and way down here, a factor of 100 smaller, is a little blip that you would never see, this is a log scale, of the resonance of the shell. I took reasonable numbers, where the gas resonance in this case is 500, the shell resonance has a Q of 100.

As we sweep the temperature, we bring the gas resonance up toward the shell resonance and we see it builds up. The gas resonance gets pulled down until it has the Q. Now they are both overlapping. The gas resonance now has the same Q, it is resonant with the shell resonance, couples energy very effectively into the shell and then radiates it off some place else. Then as you keep on raising the gas frequency you see the top mode moves up, they exchange positions. You never get a resonance at the uncoupled resonant frequency, if by resonance you mean peak.

MR. WAKELAND: What program did you use to do that?

DR. MAYNARD: You do not want to know. (Laughter)

DR. MOLDOVER: But he is good.

[Transparency]

These are the real data and what I wanted to show you is that it does not work and I do not understand it and if somebody here can help me, I would like to hear more about it.

DR. KEOLIAN: Would it have helped to put an accelerometer on the shell?

DR. MOLDOVER: What we used was a phonograph needle and we could find stuff going on. We actually did the experiment I described to you -- of course, the apparatus is long gone,

but I look back at the old data, and at the time we did it we did not know what we were doing -- and we swept the temperature through and we got a peak in the attenuation (the red curve), we got the red shift of one frequency and then the blue shift on the other side, but we got this bump and we did not see a region in the center where -- unless we got confused by some other mode and locked onto another mode. I do not understand it.

[Transparency]

A couple more perturbations. Here is our resonator, again. Transducers. The natural place, if you are interested in radial modes, is to put the transducer right in the center. Then you would see only the radial modes, because it would only couple to them, radially symmetric.

You have to know a lot about the transducer. You have to know its volume, you would to know its compliance, its frequency-dependent compliance. You have to know its surface area, because you get the largest pressure swings, the largest temperature swings, the largest heat losses, so that is not a good place to put transducers.

You would have to have wires going to it and wires are really nasty perturbations, because even very thin wires look to the gas as though they have a viscous penetration length-depth thickness. You might be handle that, but wires vibrate and I think it would kill you.

One trick was not to put the transducer where it would give you only the radial modes. This, of course, is not particularly favorable from signal-to-noise ratio, but we are not concerned with transducer noise, we are not concerned with kT noise, we are concerned with the kind of noise where the investigator does not know what he is doing. That is noise for us, it is the noise that comes after the results are out, the instruments are used and you try to understand them.

These are actually quarter-inch B&K microphones and I think that is about to scale. The microphones fit in these holders that we had to make ourselves. One of them we drive at about 100 v or so, A.C., without the bias, so that is the sound source at half the frequency we want, and the other is the B&K -- I guess we had to make our own triaxial cable. They would not sell us a triaxial hermetic feedthrough. For years they promised to make it, but they never did.

[Transparency]

Another detail -- actually this was stimulated by Bonzo's show -- was we went to a lot of trouble to put a valve on the top here. Not only is it a nuisance to make but you had to operate the valve, and then when you closed the valve there is friction involved and that heated up this thermometer here and you had to wait another hour for that thermometer to get back to where it should be, so that was a real pain.

[Transparency]

In the next version we will use a muffler. You need a hole in this resonator, because you have to get gas into it. You can calculate the admittance of a hole or a waveguide and do the frequency perturbation. If you have a hole, sooner or later it is going to be resonant with the modes in here, and when the mode in the hole is resonant with the gas resonance it is very effective in coupling energy out of the resonator.

Remember when Bob had the loudspeaker horizontal we could not hear a thing. He brought up a resonator and it coupled the energy very effectively out to us.

The trick will be to make the length of this hole a radius, and it turns out that the resonances of this waveguide will then fall between those of the radially symmetric resonances

for all radial modes. We do not need a valve, we can just close this A.C., we will have an open D.C., so it is a muffler. That is the next version.

[Transparency]

We are just about done with the acoustics part of this. We ended up working with the first 5 radial modes that varied from 2.5 to 10 kHz. We ended up going from 0 to 5 atm. We fit the speed of sound squared, or the resonance frequency, to a polynomial in pressure.

Written this way, the first term is 1, of order 1 (it is actually the quantity $1, \gamma RT/M$). The first correction term is 1000 ppm. The next correction term is 10 ppm. There is a piece that is 2 ppm here and 2 ppm here, so it is well behaved. If we had alternating sines and pieces of equal size, this would be a very lousy expansion, and that is actually how the people at NPL got themselves in trouble.

They worked at low pressures, they fit a quadratic equation to data that did not justify it and then they got the wrong intercept. Their slope was way off. This coefficient A_1 is the acoustic virial for argon and it is known from theory to be about this value. This is our measured value for the slope.

They overlap, but we have more precision than the theory, so we are actually measuring something, but that they agree tells us we are not doing something crazy such as overfitting data, which is one of the sorts of noise that the experimenter can introduce into what he is doing.

This term that goes as $1/P$ -- I did not mention it, but there is something called the thermal accommodation coefficient, which is a kinetic theory concept. If one looks closely enough, one sees the mean free path. If somebody here can provide me with a good argument that the thermal accommodation coefficient is 1 for practical purposes under these conditions, we would love to get rid of this allowance for the fact that it might not be 1, because this is badly behaved at low pressures. As you go down in pressure, the mean free path becomes larger as $1/P$ and one has to allow for incomplete accommodation.

This term was determined in another laboratory from some very high-pressure measurements. Since it goes as P^3 , it shows up in high-pressure measurements a lot better than it does in ours.

I propose to take a break 5 minutes early, because this seems like a convenient place.

[Transparency]

DR. MOLDOVER: We talked so far about resonance frequencies, measuring them and understanding them, and that is only one part of measuring the gas constant. The eigenvalue we will take from Rayleigh; he got it right. We have to measure the volume of the resonator and the molar mass of the gas used. I will say a few words about this, very few, and I will mention this because it turned out acoustics popped up in that, too.

[Transparency]

In the 6 years between 1955 and 1961, a man named Cook in England measured the density of mercury twice. He did it 2 different ways. He made himself a quartz cube, a cubical box out of quartz-lapped plates, and measured the dimensions optically and filled it with mercury and weighed it.

He also took a bucket of mercury and weighed a tungsten sphere in it. He got the same answer within a few parts in 10^7 , and for this he was made a knight. If the same guy had done the same thing today, he would have been fired for wasting time. (Laughter)

It turned out the standards labs were no dummies. They got about 500 lb of mercury in a big pile and mixed it all up so it was all the same, compared their mercury to Cook's, and now the standards labs around the world have a bucketful of mercury whose density is known as well.

Our problem was to measure the volume of the resonator by filling it with mercury and weighing it, but that is not exactly the way to do that kind of thing. What you want to do is weigh a bucket that is made for weighing in fits in balances, then weigh the bucket full of mercury, then pour the mercury into the resonator and weigh what is left over in the bucket and that difference will be how much mercury you poured in here.

One last thing, of course, is we have to know the volume at the triple point of water, not at the temperature that balances work at -- balances work at room temperature. We have to do the volume measurement and the weighing while the resonator is cold, without its microphones, of course, and then we have to correct up to room temperature.

I am not going to tell you all about it, I will just mention one question is how do you know that the resonator is full of mercury? (You cannot see inside there.) What we did is replace the valve with a fitting that included a glass tube and an expansion volume and we fed the mercury in a capillary tube through here.

We filled it with mercury up until the mercury level was in the tube like a little mercury thermometer. Then we applied pressure to that tube and the more pressure we applied, the more mercury we pushed into the resonator. You can measure the height of the column of mercury as a function of the applied pressure and you can do this going up and down, rising and falling mercury.

You should know that slope, because that depends upon the compliance of the mercury and the compliance of the resonator, neither of which is rigid. If you get more compliance than you expect, you have a bubble.

We did measure the compliance of the resonator -- acoustics came in there. We had the resonator full of some gas and we were looking at a resonance and we reduced the pressure in the surrounding vessel but not in the resonator, so there was now a pressure difference across the resonator and we could measure the change in volume by the change in resonator frequency.

We knew the compliance of the resonator -- we could also calculate it from a handbook value of steel, but if you are doing a standards measurement you cannot trust what you read. We knew the compliance of mercury (I am sure somebody got that from an acoustic measurement) and things worked out and we did get the same answer when we did it a few times.

Another comment: Why did we use argon? Physicists' favorite gas is helium, not argon. The question there is what do impurities do to the speed of sound in a gas. If we look at the impurities you are likely to run into if you do not properly evacuate your apparatus or it outgases, if you put a part per million in water in helium it will change the speed of sound 4 ppm. It will change the speed of sound in argon only 0.1 ppm. You do not have to be quite as clean if you work with argon than with helium.

The same thing for heavy impurities. Put a little xenon in helium and you get a speed-of-sound change of 30 ppm. Put it in argon and it is 3 ppm. It is just a heavier mass.

[Transparency]

There is another issue and that is argon comes in 3 isotopes, Ar-40, Ar-36, and Ar-38. One knows the atomic masses of these isotopes very well; I think those guys got it down to parts in

10^8 or 10^9 , but how much of each different isotope you have in argon that comes out of the cylinder might vary. We do not know, it is just not known well enough.

The people who do mass spectrometry cannot measure these relative abundance ratios well enough to determine that at the level we want, so we, for about \$1500, bought some liters of isotopically enriched Ar-40 (that is the most common isotope), and they were in the business of selling the other isotopes, 36 and 38, to other people, and the 40 is what they threw away. Well, we wanted the 40.

We had something that was almost entirely isotopically enriched Ar-40 and we figured we would compare the Ar-40 to the argon on our cylinder and we would get the ratio of the masses, because everything else would be the same.

[Transparency]

When we looked at the Ar-40 from the Mound Laboratory and measured the speed of sound in it, we found these results. The speed of sound had a funny frequency dependence; it went up as the frequency went up. The excess half widths, which we can measure very well, and this is excess relative to not what our theory is but relative to ordinary cylinder argon, there were extra losses. At that point I refer you to Hank's undelivered lecture here for a discussion of relaxation in gases.

We knew that carbon dioxide component of air has a slow relaxation and it turns out that that relaxation time depends upon the pressure of argon or some other gases that are around. We knew we had mostly argon. We guessed that there might be some carbon dioxide in there. Then we would expect an excess loss that is this function of this $\omega\tau$, where X here is the mol fraction of carbon dioxide, and an excess frequency mol fraction relaxation time.

The relaxation times for argon and CO_2 in this situation were known, so we could fit for the mol fraction and we came up with 35 ppm and we came up with this curve and this curve and we said, aha, the Mound Laboratory sent us argon that was isotopically enriched Ar-40, but it was chemically pretty shitty.

[Transparency]

What we then did was expose the argon to a heated titanium gedder [phonetic]. We did not have a good way of doing it, but we did the best we could. You compare the speed of sound, now in parts per million, to the stuff we were using, which we call Ar-M, for Matheson. Unprocessed, it gave this speed of sound. When we cooked it overnight or one day, the speed of sound changed. When we cooked it some more, it stopped changing.

What we assumed, and this is a little hole in the argument, is that we cooked out all the carbon dioxide, which reacts with the titanium, and had only argon left, Ar-40 left. Further cooking does not change it, so we then took this difference, the speed-of-sound difference, in parts per million between the everyday cooking argon and purified Ar-40.

We also had another cylinder in the lab that we got from Airco, made a few years later, and the nice thing is the Airco and the Matheson cylinders gave us the same speed of sound to a few tenths of a part per million. If you buy 69's argon from one of these manufacturers, you probably know the speed of sound in that to sub parts per million. I do not know how they do it, but it comes clean.

[Transparency]

I have told you about measuring the volume with mercury. I have told you about measuring isotopic abundance. I have told you about measuring frequencies and extrapolating to 0 pressure. I did not tell you about temperature measurements. What I will tell you simply is that we quit this project when the errors from each source were more or less the same at about a part per million. The reason you quit at that point is that you are going to have to fix everything if you want to do better, and that is when you need more energy and you leave it to someone else.

[Transparency]

As it happened, the results got published in the *Asahi Simbum*, which is, I am told, the Japanese equivalent of *The New York Times*. There is the gas constant R. (Laughter)

This is our work, which is about 5 times better than the earlier acoustic work. The one that got us started is off the chart; eventually they did it over, again, Quinn, Kokoff, and Chandler. Chandler retired and Kokoff and Quinn both became administrators after that, so they got out of the business and we think we did 5 times better than they did after they stopped overfitting their data.

[Transparency]

Maybe to point out a loose end or two, I told you why we used the first 5 radial modes. The sixth one happened to be close to the resonance of the shell. The seventh one falls on the shoulder, or very close by accident, of spherical Bessel functions to the 13.2 mode. This will be 13 2L plus one degenerate, a 27-fold degenerate.

We would have to pick out the resonant frequency here from there if we wanted to use the next one, if we wanted to use that mode in our work.

[Transparency]

We tried to do that and it turns out you can fit it, but there is a diagnostic that tells you that you are in danger, and this is the excess half widths of the various radial modes in parts per million.

This is argon and this is helium. The 0.7 mode I was just talking about, as you go down in pressure, the excess half width starts going up for both argon and helium, so that is an indicator that we are not resolving the modes very well and it would be not good to use that mode in extrapolating to 0 pressure.

[Transparency]

I want to say a few words now about non-radial modes. We spent all our time so far talking about the 0.2, 0.3, 0.4

-- actually, 0.2 is right next to the 3.1 here.

[Transparency]

Here is a blow up of that, 0.2 and 3.1. The reason why the 3.1 is so small is we put our 2 transducers at right angles -- I do not know if you noticed that. One was +45° and the other was -45° at vertical. The symmetry of that wave function is the same one that Jay showed us last week. It has lobes that come out on opposite sides, 180° apart, like a dipole. By putting it 90°, one transducer could excite it but the other would be at the node and would not see it, so we made this small relative to what we were looking at by choice of transducer location.

[Transparency]

Let's go back and look at that mode a little bit and see what we see. We sweep the frequency through it and we measure the in-phase and quadrature responses. Look at the quadrature here. It looks like a resonance pointing down and here it looks like there is another resonance here a little bit broader. It looks like there are only 2 modes.

Suppose you fit that to a sum of 2 complex Lorentzians. You have 2 complex amplitudes, 2 complex frequencies, and a complex background, so you are talking about 12 parameters. If you fit that, you get these deviations. It turns out the best fit gives you one mode here and one mode here -- this one is a little wide. The deviations are certainly not random; there is about 2%.

There ought to be 3 modes there. It looks like there are 2, but let's fit it with 2. Let's go for broke and put in 16 parameters; you can fit anything with 16 parameters. Not true. Anyway, you fit it with 3 complex frequencies and you find there are 2 that are close together and one far apart, and then the half widths of all 3 are the same and what they should be for the thermal and viscous damping for this mode. This is not a radial mode.

[Transparency]

In fact, the gas motion looks like this. It sloshes back and forth.

[Transparency]

Let's now look at the pressure dependence of this mode. Remember, this is what we found with the radial modes. As we went through the shell resonance, below the shell resonance we had negative slopes; right near the shell resonance there is a huge positive flow; and then at a higher frequency still it tends to come back down.

Here is the non-radial mode. Here is the doublet and the singlet, and they do not extrapolate to 0; they do have a nice linear pressure dependence, but they do not extrapolate to 0.

[Transparency]

Let's look at those slopes and that intercept. We are going to have 2 intercepts that are negative and one that is positive and then a whole bunch of these triplets.

[Transparency]

These are the 2 negatives and the one positive. What you see in the center is the average, just the arithmetic average of the intercepts, and that is very near 0. This is a nice illustration of that theorem that you can prove in perturbation theory, that the splitting of any one of these non-radial modes is first order but the average of the multiplet is second order in the deformations.

[Transparency]

Jim Mehl and I, when we sorted this out for acoustic modes and electromagnetic modes, wrote this paper that we thought was the best paper we had ever written and we got no requests for reprints whatsoever. We tried to put it in *Physical Review Letters* and it got turned down there, but that is not surprising. Then we published it in *Physical Review*. Instead we had to go do the work. That was before the fact.

[Transparency]

I told you about the intercepts. The slopes are kind of amusing, too. The slopes are very small in the center here. They diverge at low frequency now and they also diverge up here at some high frequency.

[Transparency]

The low-frequency divergence has a very interesting or amusing, to me, interpretation. The gas is sloshing back and forth in the resonator and the resonator is approximately a free body, so when the gas moves one way, the resonator moves the other, and there is a reduced mass effect.

If you look at the mass of the gas relative to that of the resonator it is linear in pressure and then you allow for the fact that you have got a sphere and you get a correction here, but there is a frequency shift that depends upon the mass of the gas relative to that of the shell. And that works out.

[Transparency]

That explains why this thing is going badly at low frequencies. At high frequencies you start getting into egg-shaped vibrations of the shell.

The reason I talked about radial modes, no one in his right mind would use a non-radial mode to do acoustics in a sphere. Why fight with fitting 16 parameters and separating out these overlapping modes? It is an utterly stupid thing to do.

However, if you want to measure the thermal expansion of that resonator using microwave resonances, you are stuck. We have a vector field. There are not radially symmetric microwave resonances.

[Transparency]

We now switched to looking at something called the TM₁₁ mode. It is a microwave resonance. We are up in the 1400 MHz region. We see a doublet where we expect to have found a triplet. We can fit it, it turns out, with a doublet very well to something better than a part in a thousand.

We could not resolve the triplet that we knew was there, but we know it is a triplet.

[Transparency]

By fooling around with the probes, what corresponds to microphones, we could tell that the lower component was really the doublet and the upper component was the singlet. We could take the weighted average of these two components to see how the volume changed with frequency.

[Transparency]

We could, therefore, measure the thermal expansion of this shell by measuring the change of the microwave resonance frequencies and we did it for three different microwave triplets from 200° to 300° or so, and you get a fit where the deviations are parts in 10^7 .

One can measure now, knowing that they are non-radial modes and knowing that they track the volume, the volume change of this resonator as one changes the temperature. The first thing we did was run it between room temperature, where Cook's mercury was calibrated, and the triple point of water, so this is changing the temperature about 30° and the volume changes some 1400 ppm.

We did it by mercury dilatometry and we did it with different microwave modes and we got the same answer, so we said, ah, we do know what we are doing.

PARTICIPANT: What is dilatometry?

DR. MOLDOVER: Dilation to expand -- you put a capillary tube into the top of the bucket and watch the mercury level go up. Physical chemists do this all the time.

[Transparency]

This graph was just to tell you that measuring the temperature at the level we wanted to also had headaches. We have our resonator in a pressure vessel in a stearic liquid bath and here is the north pole and the south pole.

Here is the temperature difference between the thermometer and the north pole and the south pole and you see that the north pole is warmer than the south pole, just as on the earth. The reason why the north pole is warmer is it is closer to the room and some of the room temperature is leaking in here. But we are talking about only tenths of a milli-Kelvin, so we can measure the temperature.

[Transparency]

Here are the results of a bunch of acoustic measurements at different temperatures, what we believe to be the right answer. What we are doing now is comparing what we believe to be the thermodynamic temperature as determined by acoustics to the scale that the wise men agreed on in Paris in 1990, so it has a subscript 90.

All the points, except the ones I have circled are from other kinds of thermometry, gas thermometry or radiation thermometry. This is not a fit to all the data, which is what they had to do; this is a fit to data we have taken recently. These blue triangles were a similar method with larger areas that was carried out in England.

This is where thermometry stands, and I will take one slide to tell you where it is going.

[Transparency]

This is the thermometry problem. Around room temperature it is no really big deal, but by the time you get to 500 K or 600 or 700 K there are two different kinds of constant volume gas thermometry (this is where they measure pressures and volumes).

There was one set of data that looks like this that was done in the 1980s and then they used the apparatus -- the guy who did this spent his whole life on it, he retired, the apparatus was taken over by somebody else -- and with the same apparatus they got this set of data.

The wise men in Paris said, "Dammit, we're going to have a temperature scale in 1990, come hell or high water," so they split the difference. (Laughter)

That is not the way to do science. You do not do that. Yes, sir?

DR. HOFER: The units on the left are in milli-Kelvin?

DR. MOLDOVER: These are milli-Kelvin, yes. These are the differences between -- well, there are two different plots. The last one was the difference between the temperature that we determined and what the 1990 scale is. And this is the difference between the 1990 scale and the 1968 scale. The last time the wise men got together was in 1968.

PARTICIPANT: Those are big numbers.

DR. MOLDOVER: You think it is easy? (Laughter)

Come visit us.

MR. WAKELAND: The 1968 data is the lower set of data there?

DR. MOLDOVER: The 1968 scale is the base line here. Then this was done in the late 1970s, early 1980s, and this was done in the middle 1980s with the same apparatus.

PARTICIPANT: But your blue markings are nonlinear.

DR. MOLDOVER: You can switch between temperature and ppm in the speed of sound. I am sorry, thank you for asking. Bad notation, very bad notation.

The difference between these two gas thermometry measurements is equivalent to 10 ppm in the speed of sound. The difference between these two sets is equivalent to 25 ppm in the speed of sound. We think we can measure the speed of sound to a few parts per million, so we think we can settle this, and we are working toward it.

[Transparency]

I want to mention only one idea that I picked up from hanging around chemical engineers a lot. They do wonderful things with flowing gas systems; that is how they get pure gases. They do not leave them sitting around long enough to become contaminated. They flow through a plant, distill them, they are always working in flowing systems, so we are going to have the capability of flowing gas through here and that is why I am interested in mufflers, so the resonance time of the gas in the apparatus, instead of being days will be minutes.

I think I want to put this aside and get on to more practical, less aesthetic, applications of acoustics.

[Transparency]

We have gamma in all of our equations, so that tells us speed of sound depends upon the heat capacity of a gas. We are going to run that backward; that is, we are going to measure the speed of sound really well and get out the heat capacity, so you guys can know what to put in your equations.

The other thing that one really wants is the equation of state of a gas, which means I want a way, from measuring the temperature and pressure, for which I have sensors, to know what the density of the gas is. That means I want to know these temperature-dependent corrections.

I am going to be talking only about gases, and by gases I mean something less than about, let's say, the critical density of a fluid. Critical density would be a third the liquid density. You do not want to use resonance techniques for liquids, because the resonator shakes and becomes part of the problem. There are plenty of good pulse techniques and so forth.

[Transparency]

The first thing we did -- here is something that did not work well. We built a small version of that volleyball-sized resonator. This one you can hold in your hand, it is about an eighth of a liter, a cupful. It is a little resonator in a little pressure vessel and you set it up and start measuring the speed of sound in alternative refrigerants like our 134A, and you can measure the speed of sound going from 200 to 340 K as a function of pressure and you get all this information.

[Transparency]

It worked really fine until we got to a compound that was called (the third one we tried) E134, and this is the compound -- do not ask me what it is. What we found is the speed of sound depended upon time between the time you put the gas in the resonator and you made the measurement. You get this speed of sound in 150 ppm, this is overnight, and it just kept on going.

That forced us to think about what we were doing, being naive physicists.

[Transparency]

What we had done foolishly is put transducers in the pressure vessel and wires in the pressure vessel, not in the resonator but in the pressure vessel, and O rings, and you had a gas

mixture -- it was not intended to be a mixture but it is a mixture -- and one component is more soluble in the O ring than the other, so the composition was changing right before your eyes.

So we said, all right, O rings you will never lick. If you ever want to do an interesting experiment, take an O ring and expose it to carbon dioxide at high pressure. Then take your vessel apart and you can take your O ring out and watch the O ring swell up before your eyes and then shred itself.

Gases dissolve in polymers and some of them dissolve really well.

[Transparency]

We decided we needed an all-metal apparatus and as a prototype we made a cylindrical resonator, just because cylinders are easier to make than spheres. We had done these microwave experiments and one does not put a microwave generator inside a thermostat. One has a microwave on a relay rack and one brings the microwaves into the thermostat using coaxial cables.

We asked why can't you do that for sound? We have an acoustic source, acoustic waveguide, and bring it into the resonator and look for resonances, bring it out to a detector. Here is the trick we used, which might not be necessary, but it did not hurt.

This is the end of the waveguide and there is a diaphragm welded across here. There is a little screen. You buy these tapered tubes from B&K and they have screens in there and they kill reflections in the tube at this point, because we did not want standing waves in the tube to go back and confuse us with the resonances in the resonator.

The guys who do microwaves have 50 ohm terminations in all their coaxial cables, so it is the same story.

[Transparency]

You pay a price for this. If you have a pressure difference across that diaphragm, it gets stressed and the transmission goes way down. If you stress it too much, it pops. What you have to do is balance the pressure inside and outside, and you can do that.

[Transparency]

Just as with the spherical resonator you can measure the speed of sound with a lot of different modes. There is a radial mode and there are modes that run around the outside of the resonator, like this one, and there are longitudinal modes. We got consistent results at a certain level.

[Transparency]

Here is what the apparatus looked like. This gives you an idea of the scale. You can hold it in your hand easily enough. This is a waveguide coming down and there is a brazed joint over there. This is a fill tube. A radiator clamp holds a thermometer on to this.

MR. WAKELAND: Could you say a little bit more about that screen? The idea it has thermoviscous dissipation that acts like a real --

DR. MOLDOVER: It is a resistor, that is all.

MR. WAKELAND: Those tubes that are coming in look pretty small.

DR. MOLDOVER: They were tapered, yes. We wanted a fairly high impedance -- we thought we did -- at the diaphragm. I am not sure that I understand this fully.

MR. WAKELAND: When you say screen, you mean something very, very fine.

DR. MOLDOVER: Yes. We bought it from B&K. These are horns that are used with probe microphones. If you want to measure the sound in a furnace, you stick this little tube into the furnace -- maybe somebody here can answer better than I can. I am not an acoustician.

DR. SABATIER: Something like Bonzo's Kleenex that he was putting in the end of that rectangular trough.

DR. KEOLIAN: There are some tricks.

DR. SABATIER: If you want to make a point source, in audiofrequencies they often drive it with a long tube that necks down, and if you do not put a screen in the end of it and you look at your temperature response, you get big sharp peaks. If you put the screen there, it dampens it down and makes a smoother transmission course. The impedance of the atmosphere is better matched.

DR. MOLDOVER: So that was the resonator.

[Transparency]

Here it is installed in a bath that you stir with oil or fill it with alcohol or oil. There is Keith Gillis, who got this working. He is a great guy and a lot of the work I am telling you about just would not have gotten done if he were not around. A lot of people here know him, I think he was at one of the schools before.

That was the first practical apparatus. We did a lot of work on it with alternative refrigerants. I think I may have given you a list. We never bothered to put a sphere in there. We did not need one, it worked well enough as a cylinder.

[Transparency]

Here is the next version with John Hurley. It is more or less the same apparatus that Keith built, except it is inside this cabinet. You see here is a nice panel that shows where the valves are. We are really getting into engineering here.

[Transparency]

When you see something like this, you know that this is being set up for routine work. Yes, sir?

DR. HOFER: Are those electrically operated?

DR. MOLDOVER: They are all air-operated.

[Transparency]

This really is a great way to measure thermodynamic properties. What you do is you fill this resonator, whatever you make it, cylinders, spheres, with argon, and now you know how to calculate the speed of sound of argon, it is $\gamma RT/M$. You know the argon that comes out of a 69 cylinder really is 69, because we checked that out. You know T as well as you can measure it.

It turns out that you do not need to measure T very well. The reason is, if you put argon in there in your calibration and you have got the temperature wrong, and it is wrong the same way when you put in your test gas, you are measuring the ratio of the speed of sound of argon to the test gas.

If you look at propane, the speed of sound changes 100%, going from 210 to 460, but with argon the ratio of argon to propane changes only 16%, so you do not even have to know the temperature that well.

If you are going to calculate the heat capacity, or gamma, from your speed-of-sound measurements, take my word for it, the ratio of the heat capacity to R is Mc^2 , the speed of sound, Mc^2-RT . We are not going to measure R and T . We are going to put argon in there and measure the resonance frequency of argon. That is our calibration.

Let's just put it in here and what we see is the heat capacity at some temperature is 1 divided by the frequency ratio at that same temperature and the mass ratio at the same temperature, so you do not even have to know the temperature very well. Not only that, you can have a temperature gradient on the resonator from top to bottom, so long as you have the same one as when you do your test gas and your calibration.

Everything cancels out. This is wonderful. The only thing that does not cancel out is this mass ratio but argon comes really pure, so you had better be sure that your test gas, if it is propane, is really propane, because this denominator tends toward 1 as gamma goes toward 1. Another way to write it would be to write this as $1-1/\gamma$.

If you want to get 1/10 of a percent in heat capacity, you had better have your mass to 1/100 of a percent for something like propane. I think this is the standard now for measuring heat capacities.

[Transparency]

We have done a whole bunch of fluids, alternative refrigerants that Steve mentioned. The underlined ones are sponsored by the Navy for its particular purposes. We have done some work with helium-xenon mixtures. We are now getting into the gases used in the semiconductor processing industry. We have done boron trichloride, we have done chlorine with no trouble at all. Hydrogen bromide is giving us headaches. It decomposed in the supply cylinder and we purified it.

I might mention why we are interested in these gases.

[Transparency]

Here is a device called the mass-flow controller. I do not know if any of you is familiar with it. It is used in the semiconductor industry to feed gases into processes. In a semiconductor you put in a fluoride onto a heated substrate and the fluorine gets deposited. They flow gases in, you want to dope your semiconductors. There is a meter that is fed to a valve. The meter is a bypass and basically a fancy hot-wire anemometer. There is a heater and thermometers in here.

There is a whole bunch of different proprietary ways of doing this. I think one of the cleverest ones is you modulate the temperature here and you detect how long it takes to get the modulation to flow downstream. I think that is pretty tricky. There are other schemes.

In order to understand the meter, you have to understand the pressure-density-temperature relation for the gas, you have to know its heat capacity, its thermal conductivity, and its viscosity.

If you use devices like this to prepare a gas mixture in a pilot plant you want it to work on full scale, so your meter needs a real calibration, not just a fake number. That is an example where you need these properties.

[Transparency]

The gases they are talking about are a long list of "nasties." The people who make the meters do not want to calibrate them with these gases. We do not want to calibrate them with these gases. They are very happy to calibrate them with nice things like the chlorofluorocarbons,

they will do that for you. What they have are calibration factors that they have obtained by a cookbook, and what we are going to do is give them the virial coefficient of viscosity and the Prandtl number, so that those calibration factors are rationally based rather than based on best guesses.

So do them once; that is the motivation.

[Transparency]

I may have handed you this graph in the notes. The point here is that here are 2 isotherms at 300° taken with a sphere and there is one taken with a cylinder and they agree. The difference here, which is about 100 ppm, is a difference in purity and not a difference in our ability to measure speed of sound. Our limiting factor in all this work has now become gas purity. If you use it intelligently, you can use any old resonator.

[Transparency]

The next question that comes up is what do you do with the data. We now have these automated setups, and you have seen it. Here is a refrigerant. We have gone from 240 to 400 K, lots of speed-of-sound data. You can fit it with polynomials and it fits in parts to 10^4 . RMS deviation was 2×10^{-5} . Impurities cause larger systematic errors. And this tells you what the surface looked like.

[Transparency]

From those data you can extrapolate it to 0 pressure -- sorry, I meant to point that out -- when you fit it, and that will give you the ideal gas heat capacity.

Here is the ideal heat capacity for the same fluid and you see it is 10R going up to 12R as a function of temperature. These are deviations from the fit, 2/10 of a percent, so the precision is certainly better than that. For these kinds of molecules it becomes hard, but they try to estimate the heat capacity from these infrared spectra, find out the vibrational modes, rotational modes.

These are independent measurements or spectroscopic estimates. My own belief is that the acoustics in this case is better than the spectroscopy. That is the 0 pressure limit.

[Transparency]

We are measuring speed of sound as a function of temperature with coefficients that are temperature-dependent

-- I called them beta, gamma, delta, epsilon -- temperature-dependent functions of pressure.

Let's look at beta, the first pressure correction. It is related to the density of the virial coefficient, which is what you need to get the gas density by a differential equation. We are measuring beta, so we have this side of the differential equation. What we want is B, so we can put it here and calculate the density.

The problem for many of these fluids is there are no initial conditions for this differential equation. We need a first derivative and second derivative. That means we need values of B at 2 temperatures to have initial conditions to do that integration.

People have traditionally tried arbitrary functions and I think our contribution was to use a square well. Let's assume that a molecule has a hard core, an attractive region, and nothing out here. One can calculate for this potential from statistical mechanics. It has 3 parameters. It has the size of the hard core of the molecule, the depth of the well, and the size of the attractive region. Those parameters appear here in b_0 , in epsilon, and in this R, which is a ratio here. Anyway, there are 3 parameters, that is all I want to tell you.

You can fit the data with this guess and it turns out it works quite well. For this test case we got gas densities -- here is the equation of state. This is not an ideal gas, this is deviating from an ideal gas by 10 or 15% in density, but we are getting the density from this guess, this square-well approximation, to within a fraction of a tenth of a percent. So this is a procedure that works.

[Transparency]

Here is a sort of dimensionless plot on which you can put down all fluids, if you wish, except may helium and liquid metals. Here is the vapor pressure curve, pressure divided by the vapor pressure, temperature divided by the critical temperature. We are talking about gases.

The data I just showed you span this range. From speed-of-sound measurements in there you can get out densities, as I showed you.

[Transparency]

I want to refer you to a colleague's work that covered the whole damned vapor range. What we had to do, as we went to higher and higher virials, we had to add more and more square wells and more and more potentials. I am not explaining it clearly.

The square well associated with the gamma was not the same as the square well associated with beta, so there were 3 parameters here, 3 more here, and this is not physics, it is curve-fitting. There really is only one potential; it is just not square wells. Square wells were a good first cut, but not good enough.

[Transparency]

Martin Trusler at Imperial College took a really good 4-parameter and a 3-body term, so 5 parameters. You cannot do things in closed form any more, you have to integrate these potentials weighted by the Boltzmann factor and then fit them, but he could, with 5 parameters, get out good gas densities over the whole range.

He has persuaded me that a good physically based model for the potential plus speed-of-sound data can get you gas densities to about 1/100 of a percent and that is, I think, about the state of the art in this. We are going to have better gas densities than you could get any other way. Let's take a break.

DR. MOLDOVER: A question came up at the break and I apologize for not having given you this road map. What we are measuring in these so-called industrially important gases is the speed of sound as a function of pressure and temperature. We have the data and we use them to fit parameters in a model potential, where this is a model intermolecular potential.

Real molecules are very complicated, but once you have a model potential you can grind and get out an equation of state. The thing we want for designing, for flow-metering, whatever, how turbines work, we eventually want the density. We cannot go right from here to here, because we do not have the initial conditions for the equation of state, so assuming a potential is implicitly assuming initial conditions for this integration.

[Transparency]

New subject: To first order the speed of sound measures the temperature -- we have been talking about that in the first hour. To first order the viscosity or the thermoconductivity measures the size of a molecule.

[Transparency]

If you wanted to measure the size of a molecule and you take a physics course, you think of doing a scattering experiment and measuring the size of the shadow. When you measure the viscosity, nature does the scattering experiment for you, your molecules are scattering randomly and you measure a mean free path.

It turns out that the transport properties can usually be measured a little bit more accurately and a little bit more conveniently than this experiment can be done.

[Transparency]

It is a fun way to think about what you are doing when you are measuring transport properties, which is otherwise pretty boring. Another fun thing to think about is calculating transport properties. Using Jay's quantum mechanics, one can calculate the transport of, say, the viscosity of helium, and we have done that.

There are a lot of steps to that problem. One solves the Shrodinger equation for 2 electrons, and there are people who are experts in that. Then one gets an effective potential. What we did was we used the effective potential and did all these partial waves and stuff to get the viscosity.

This straight line is the result from quantum mechanics, viscosity going from about 10° to 1000° . The deviations from that line are the data. All the deviations are linked to one particular measurement by now deceased but excellent mechanical engineer at Brown University, Joe Keston. He had an oscillating disk viscometer that was an absolute instrument and he got this result and then he got this temperature-dependence (these circles). Someone else has gotten a different temperature-dependence tied to his, and all these low-temperature measurements are ultimately tied to his. The situation is such now that the ab initio calculation for helium is as good or, I think, better than the best measurements.

One last thing that is sort of fun to notice is that the kinetic theory results for the viscosity have a diameter, hard-sphere diameter for hard spheres. That is this line, horizontal line. Helium is definitely not a hard sphere on this scale. As you might expect, the higher you go in temperature, the smaller the sphere diameter needed to fit the viscosity, so in order to bring the red line up to the solid line, you have to reduce the sphere diameter.

The higher the temperature, the harder the collisions among the helium atoms, the smaller the atoms look, the more they overlap. At low temperatures, helium looks bigger. Yes, sir?

MR. STOREY: The solid line is like your fancy calculation and the red line is just using the hard-sphere formula?

DR. MOLDOVER: Right.

MR. STOREY: What does it look like if you use a fancier potential like a Leonard-Jones potential? Do they match better?

DR. MOLDOVER: You will get closer and closer to the solid line. If you put in the right Leonard-Jones parameters you will get very close to the solid line. That is a good question.

[Transparency]

You can do the same thing for the other transport properties in the virial coefficient helium. The one thing I wanted to mention is if you do it, we think the Prandtl number for helium is not quite $2/3$ (everybody says it is $2/3$). It differs by about $3/10$ or $4/10\%$ from two-thirds.

It would be kind of fun to see if you could measure that and I think --

DR. KEOLIAN: Isn't it a finite temperature sort of thing?

DR. MOLDOVER: It is temperature-dependent. There is nothing magic about $2/3$.

DR. KEOLIAN: Isn't it an absolute thing from kinetic theory of ideal gas, ideal monatomic gas?

DR. MOLDOVER: I can tell you what I did, which is calculate the transport properties and compute the Prandtl number. It may be for hard spheres it is $2/3$, but helium is not a hard sphere.

MR. WAKELAND: It is $2/3$ if there are no internal degrees of freedom that are excited, right?

DR. MOLDOVER: And if it is a hard sphere.

[Transparency]

How do you measure transport properties acoustically? If you look at a resonator -- this is a picture that Keith gave me -- and look at a standing wave like the fundamental mode, in the center here, where the gas velocity is high, we got lots of viscous losses. At the ends, where the gas velocity is low, the pressure swing is large, the temperature swing is large, and we have lots of thermal losses.

If we want to measure viscosity, we want a resonator where we got lots of viscous losses and very small thermal losses. Mo Greenspan, I guess in the late 1940s, conceived of a viscometer that looks like this, where you take this resonator and squeeze down the center and make the high-velocity region have even higher velocities, so you get even more viscous losses and, at the same time, open up the ends so the pressure swing is smaller and the thermal losses are smaller.

Another way of thinking about it is you take this section of the wavelength and put it here and the ends and put them here. The pressure drop is almost all across the tube and this would be a good viscometer.

[Transparency]

He called it a double Helmholtz resonator. We like to call it a Greenspan viscometer, and I hope the name will stick. He was really a great man. If you balance the kinetic and potential energy you will get something like this. We have the speed of sound, the area, the length of the duct, and the reciprocals of the volumes of the 2 chambers.

That back-of-the-envelope calculation I gave you before is very nice to apply to this device. Let's talk about the duct.

[Transparency]

Now I show it to you in cross-section. During the half of the cycle where the gas is sloshing from one chamber to the other, the energy is kinetic and it is all a half MV^2 and it fills almost all the volume of the resonator except for the viscous penetration length.

That back-of-the-envelope calculation says that the Q is going to be the energy stored, which will be the radius squared, divided by the energy lost, which will be the circumference times the viscous penetration length. That is this term here.

You can do the same argument for the thermal effect in the ends and you find that the volume -- you take a thermal penetration length from the volume, multiply it by the surface area here, and divide it by the total volume. This is the volume in the boundary layer. This is the volume in the chamber. You can see from the shape that the volume in the boundary layer as a

fraction of the total volume is quite small; so the thermal correction is small, the viscous correction is large.

This is what you want for a viscometer, so we deform that cylindrical resonator into a viscometer.

[Transparency]

There are some details to be worked out. This is a plot of pressure contours. There is diverging flow from the nozzle and there is kinetic energy in that flow, so that has to be accounted for. There is also some viscous drag around this corner.

Those of you who know electrostatics know that when you have a sharp corner like that, you have large electric field lines and the acoustic analogy is you have a large pressure gradient and the velocity, when it flows around the outside corner, becomes infinite.

[Transparency]

That, too, can be dealt with. One can round the corner -- it is an integral of infinity. Although the velocity is infinite going around the corner, the dissipation is finite.

[Transparency]

Another item that you would want to model is the effect of changing the insertion of the tube into the chamber. This tube could stick all the way out or it could be flush with the wall here and it would be like a baffle. One can compute the added resistance of that orifice as a function of insertion.

Here we have a baffle result. When the tube comes close to the far wall, the resistance starts going up and the inertia starts going up. You see from this sort of design, thinking that the corrections for the orifice are on the order of an extra radius in length of the tube, so to get a good viscometer you want a tube that is fairly long, and it is easy to get one that is 10, 20, or 30 times the diameter. There is a correction at the end that is on the order of a diameter; it depends on corners, it depends on how far away the baffles are, but if you make the corrections small and then make it -- yes, sir?

MR. SMITH: On your scale there on your graph, that is variations in radius over viscous penetration depth, the bottom one, I am guessing?

DR. MOLDOVER: Oh, you are worried about the label here?

MR. SMITH: Yes. It is picky, I know, but --

DR. MOLDOVER: No, that is fine. I would guess that that is the resistive length, the effective length, additional length in resistance compared to the viscous penetration length. That is sort of reasonable. This tube, so far as losses are concerned, you have an extra something like a viscous penetration length flowing around the corner. The inertial part is going to be something like a tube radius. I think they are both right. RD is the radius of the duct in Jim Mehl's notation. The inertial part has to do with kinetic energy.

Good question. Every question is good.

[Transparency]

This looks like a 4-cm scale. For some strange reason we put extra transducers in this one - do not ask me why. We do a lot of things that we would rather not remember.

[Transparency]

Here was an early fit to a resonance. It is a Helmholtz resonator, so this is, I believe, in air. It has got a low Q . Now you see we are going from 160 Hz to 220 Hz, so a half width is

something like 20 out of 200, so we are talking about a Q of 5 or 10. That is what you want for a viscometer, to have a low Q.

[Transparency]

What you are really measuring is the viscous penetration length, or the viscous diffusivity, or η/ρ . If you study this as a function of pressure, you get something that varies over large orders of magnitude, because ρ is varying a lot and you will never get viscosities better than you know the density.

You have to assume we have used the speed of sound to get you really good density values with that chain that we talked about, and then we are going to go get the viscosity. For helium, argon, and propane the viscous diffusivity varies by a couple of orders of magnitude, but mostly because the density is varying. The viscosity does hardly anything.

This is at room temperature. These are deviations from the data in the literature. In this early prototype we are talking about 1%-type stuff. It turned out the average was in the right place.

[Transparency]

We have this little kit of viscometers where you have the center sections and end sections bolted on and you could bolt on a duct, big diameter or a smaller diameter, or a long duct or a shorter duct, so we had a kit of viscometers. We measured the dimensions of these things with a coordinate measuring machine, which is a milling machine fitted out with dial indicators in an air-conditioned room and run by people who claim to know what they are doing.

There were no adjustable parameters. We put in everything we knew, so this is now an absolute viscometer, no calibration. We got kind of 1%-type results, sort of an indication of problems at low pressures in this one. The red one is this one -- it is color-coded. It did not follow that the one that looked most like a viscometer, namely, this one, gave necessarily the best results.

[Transparency]

But we decided to go ahead and make a proper engineered version in an apparatus with a pressure vessel and pressure control, and this has given us a lot of grief. Maybe the acousticians in the audience can tell from looking at this that there is going to be grief here.

There are PZT transducers here behind diaphragms -- again, we want to be able to handle nasty gases and we want our transducers out of the gas. We had a small gap here so that when you change the pressure of this gas the temperature changes and it would equilibrate in a reasonable length of time.

We hung it on wires to isolate it from the stirring in the bath, which is a bit of a problem, because we are at much lower frequencies now and general shaking is a bigger problem.

Everything I told you led to problems. The transducer has a bleed or a vent hole, and that is a small Helmholtz resonator. The wires are tense strings and they are resonators.

[Transparency]

The last bit of grief that I think we sorted out on Thursday and Friday of last week I am showing you here. There are 2 traces. One is with the resonator evacuated and the can evacuated. The upper trace is filled with argon at 100 kPa. Here is the Helmholtz resonance. All this stuff is cross-talk. Here is a shoulder over here.

We switched the gas to helium and we found the same grief. Not only that, we did it at 2 different temperatures and the grief moved from here to here, so grief is also gas. It changes with temperature. I think what that is, is a double-Helmholtz resonance, where sound from the transducer and the vent hole, after we enlarged the vent hole, pressurizes this chamber, runs through the duct between the viscometer and the pressure vessel and into this chamber, so we have resonance in the outside here.

MR. SMITH: Why wouldn't you cancel that out in the transmission line or something?

DR. MOLDOVER: I think it is a lot easier to just open up this gap and use some battery cable to thermally link

-- you know, we thought we would be clever and it bit us.

This is sort of fun. When you see a peak like that, we can now fit it to a part in 10^4 with good lock-in and only a few minutes averaging right around the peak. That is really neat, because we will not be noise-limited. Yes, sir?

DR. HOFER: In this apparatus I noticed the neck tube is mostly into the end volumes, whereas in the previous one it protruded only slightly. Is there a reason for those kinds of changes?

DR. MOLDOVER: It is a little thing. You do not want it near either end where the curve was flat, at the end where the curve of the end corrections was flat.

What we are heading toward is eventually we are going to use a very precise cylinder in there that is made -- the guys who do pressure measurements use pistons and cylinders. Those cylinders are made very well, much better than any old shop could make, because they really want to know what the pressure is.

Eventually there is going to be one of these things where we put in a cylinder that is really well known and that is going to be a viscosity standard if we end up understanding this thing as well as we understand the spherical resonator. We are heading in the direction of trying to understand this as well as we understand the spherical resonator.

[Transparency]

Let me switch once more to the question of how do you make a thermal conductivity device acoustically, how do we measure thermal conductivity acoustically.

Here is our cylindrical resonator, again, and here is our standing wave. We have a pressure and temperature peak at the ends, velocity peaks in the middle. We put a stack of plates in the middle here, so this is like an acoustic viscometer. We are going to have lots of damping of the sound waves due to viscosity.

[Transparency]

Let me switch to another mode, the first overtone. That has a velocity node here, so there is no damping due to viscosity. It has got a pressure and temperature anti-node there, so all the damping is going to be from heat flowing in and out of these plates.

Another way to think of this is as a totally useless thermoacoustic refrigerator where you put the stack in the wrong place. (Laughter)

We kluged together a prototype. These were some high-vacuum flanges that were lying around. For stack we used a honeycomb that Ray Radabout gave us.

[Transparency]

There is some process that makes steel honeycombs that they use for microwave shielding, letting air go through but stops microwaves. I think some people use them as stack elements. It is really ugly-looking geometry.

[Transparency]

You measure the response and you see that here is low damping and it is redshifted. This is the viscosity damped. Here it is high damping, second mode, high Q and blueshifted, redshifted, blueshifted, which is what you expect from the obstruction of some of the space inside -- Steve Garrett mentioned this.

Exactly the same area is doing the viscous and thermal damping, so if we want to measure the ratio of the thermal-to-viscous losses, the area should cancel out and we do not have to know the area really well. That is why you have a hope of making this work quite well.

[Transparency]

These are the excess half-widths of those modes expressed as a function of the penetration depth, which could be either thermal or viscous, as compared to the honeycomb size. I think we modeled the honeycomb by a circle. You could do all the stuff, triangles, flat plates, and so forth, but as we learned in thinking about the earth, you do not win a lot.

I guess you will have to look at the text version, but the top curve here is the thermal mode with propane (P.R. means propane). This is a viscous mode with propane on the bottom here. The argon results are in yellow and the helium results are in blue.

There are several things to see. As you get to smaller penetration depths, the excess width grows, and that means the effect of area grows. That does not surprise you. You are starting to look at this messy honeycomb on a finer and finer scale, so the area should grow. This curling up here does not tell you anything much.

What you can then do from these data is compute the ratio of Prandtl number -- if there were no excesses, the Prandtl number would be perfect, but we do not need to have that condition, 0 excess; if we have the same excess half width, we still get a reasonable Prandtl number.

[Transparency]

Here is what we got, what is plotted here. This is deviations in percent of the Prandtl number from what we thought it should be as a function of penetration depth. We have got results for propane, which has a Prandtl number of 0.75 for helium, and for argon. There is a definite trend with penetration depth, but had we worked at one penetration depth and used helium as a calibrant, we could have estimated the Prandtl number of propane or argon. I would say this is a demonstration; I think we have a long way to go. This has not been engineered.

[Transparency]

I would think you would want to do it over again with a cleaner geometry. The next version has been assembled but not tried. One of the things I worry about is the screen being fastened really firmly in place; we do not want any stray losses.

MR. WAKELAND: Why do you want to use a screen?

DR. MOLDOVER: You could use anything. We want the area that we switch from thermal to viscous to be large compared to the area of the resonator and we want a reasonable Q.

MR. WAKELAND: I was just curious. I have never seen that fancy hexagonal stuff before.

DR. MOLDOVER: It is available. It will take high temperatures and stuff. It is stainless steel and nasty materials.

[Transparency]

I want to say a few words about mixtures and then I am going to quit. The thermoacoustic refrigerators -- a pin is the optimum geometry for a stack, I am told, if you can only figure out a way to suspend a bunch of pins cheaply and rigidly. The viscous penetration depth measured out from the pin, the dissipation there is wasted, the energy dissipated due to viscosity in that region is wasted.

The energy that is pumped up or down the stack is the energy within the thermal penetration depth of the solid. The ratio of these 2 areas is the Prandtl number. The Prandtl number, then, is a figure of merit and has something to do with how efficient such a stack would be and you want a small Prandtl number.

If you look at the noble gas mixtures, the smallest Prandtl number comes for helium-xenon mixtures of this composition. Steve emphasized yesterday that small Prandtl number is not the only thing you want, there are other things you want to optimize in a design. One thing that people initially complained about was the cost of xenon (it is \$13.50 a liter), but this mixture, which is very close to the optimum one, is only \$1.35 a liter. If you are going to spend man hours or man days on tests, you are spending hundreds of dollars on it. You might invest a few hundred dollars in a fluid, if it would help make the machine more efficient and not sacrifice something else.

[Transparency]

In the same program that I told you about we are doing it with mixtures and we have done a fair bit of work for helium-xenon mixtures in this temperature range, from 200 to about 400, and this pressure range, that is, 0 to about 1.5 MPa, we have measured speed of sound. The notch cut out of the data is because we are close to where xenon liquefies.

We have also measured the equation of state up here, so we really do in this case have the initial conditions to integrate. We have a very accurate model of equation of state under these conditions.

[Transparency]

When we did the speed-of-sound measurements we were using the cylindrical resonator I showed you and we published a paper with these deviation plots. This is helium, that is xenon, and these are mixtures. This is actually the ab initio base line, so we have allowed for the pressure dependence of the speed of sound in helium from the quantum mechanical calculations. No parameters here. Xenon has some parameters and the mixtures have some parameters.

These curves are cut off at 4 atm. I have not shown you the deviations below that. That was done deliberately. We did not want to confuse the people reading the paper, but I am quite willing to share with you what is hidden.

[Transparency]

This was the xenon-rich mixture. This is a helium-rich mixture. Below about 4 atm we are getting speeds of sound diving down here and generally becoming a mess. We knew speed of sound did not do that, so we had to look more closely at our apparatus.

[Transparency]

We found the speed of sound was time-dependent. We do not have rubber O rings any more, we have learned that lesson. When the resonator was at 350 K and 10 atm, the speed of sound went up a little bit. When the pressure was down to 2 atm, it went up 5 times faster, so we are looking at a time dependence.

When the resonator was dropped in temperature below ambient, the speed of sound went the other way. We are looking at a time dependence.

[Transparency]

This mess, which looks like a pressure dependence, is not at all a pressure dependence, it is a time dependence.

[Transparency]

What is happening is shown schematically. Here is the resonator in the oil bath and there is, up on top, a pressure transducer that is somewhat above room temperature (it is kept warm to prevent condensation). The resonator might be either warmer or colder than that pressure transducer.

What happens is the xenon preferentially separates in the temperature gradient and flows to the cold place, so the gas de-mixes. This is a well-known effect called thermodiffusion. People who handle gases that are stored in tanks mix gases. They will actually take a tank and roll it.

If that happens in the course of hours over the space of many centimeters in a lab apparatus, it can also happen during each acoustic cycle.

[Transparency]

There are 2 processes that this calls to mind. This is an expression for the attenuation in the bulk of a sound wave propagating in free space. There is a viscous contribution, there is $\Gamma - 1$ times the thermal conductivity. This is the thermal contribution.

There is a diffusion coefficient contribution that has 2 terms, a term involving γ , again, the heat capacity ratio, and the thermodiffusion ratio, a static property, an equilibrium property. What this is telling us is that in the course of an acoustic cycle a temperature difference is set across the sound wave and that xenon and helium will tend to separate, with the xenon flowing toward the cold place and the helium flowing to the hot part of the wave. How much flows depends on the diffusion coefficient.

There is a second term here, so even if γ were 1 and we had a mixture of fluids of 2 different masses, so there was no longer a thermal effect, there would still be an extra loss mechanism and, as near as I understand it, it has to do with the following idea.

Think of a gas and let's paint every other molecule red, or make it 100 times as heavy as the first, so we have a big mass difference, red and blue molecules, heavy and light ones. When the light ones collide with each other, they exchange energy and momentum very effectively. When the heavy ones collide with each other, they exchange energy and momentum very effectively. When the light ones collide with the heavy ones, they exchange momentum quite well but they do not exchange energy at all; they just bounce back off.

This is really important in plasma physics, where the light ones are electrons and the heavy ones are ions and you have 1000:1 mass ratio, but it is also important in helium-xenon mixtures.

MR. WAKELAND: You mean there is an M1 and an M2 and then --

DR. MOLDOVER: Mass of the helium and mass of xenon.

MR. WAKELAND: So M, the un-subscripted one?

DR. MOLDOVER: It should be the mass of the lighter one.

MR. STOREY: Shouldn't there be a mass diffusion term floating around, just regular fixed diffusion?

DR. MOLDOVER: The 1-2 is a diffusion coefficient.

MR. STOREY: -- those two coefficients there should be an additive term that is proportional to the concentration gradient? Or do you just combine that with that Gamma-1 term?

DR. MOLDOVER: I have copied it from Kohler.

MR. STOREY: The normal mass diffusion vector should have more terms -- there should be a term that is a diffusive-type term. Then there is a pressure gradient, a temperature gradient, and then also a gravity term.

DR. MOLDOVER: I defer to the expert. Hank, do you want to say anything about that?

DR. BASS: I am not sure where the gravity term enters in. It does not enter into these equations. This is the correct expression for the diffusion term.

MR. STOREY: I agree on the 2 terms, but there should be a fixed diffusion term, right?

DR. BASS: There might be an additional gravity term, but gravity is not included in these equations, it is not included in the force of these equations.

MR. STOREY: Right, that is probably neglectible, but shouldn't there be, like viscosity or thermal conductivity, the mass diffusion?

DR. BASS: I am not sure. First term, viscosity, second term, thermal conductivity, third term is mass diffusion.

MR. STOREY: No, there are 2 mass separations but not a mass diffusion, like a fixed diffusion.

DR. MOLDOVER: But you need a gradient to separate the masses where you have a gas at equilibrium, and it can be separated by pressure -- you mentioned that.

MR. STOREY: Right. The temperature and the pressure will want to separate and the normal mass diffusion will want to bring it back to --

DR. MOLDOVER: We already have a mixed gas. You do not have an attenuation process unless the sound wave somehow puts a gradient in the quantity you want. The sound wave does not create a concentration gradient except by creating pressure or temperature gradients.

[Transparency]

I just want to show you that for helium-xenon mixtures the diffusive part of the loss is larger than the thermal or viscous part of the loss. This is not likely to influence refrigerator operation, because we are dealing with losses other than bulk losses, but if you want to set out to measure the properties and understand them, I think it is good to know this exists.

[Transparency]

In one last slide here I want to acknowledge that over the years I have had the good fortune to work with many collaborators, postdocs, guest researchers visiting from other countries, people on sabbatical, and there is no way that any of this could have been accomplished without any of these people.

SCANNING ACOUSTIC MICROSCOPY: LENSES, TIPS & SONOELECTRONICS

Calvin F. Quate
Ginzton Laboratory
Stanford University

[Slide 1]

DR. QUATE: Good morning. I want to thank Logan Hargrove for inviting me. Hank Bass, thank you for arranging all of this. Libby, thank you for helping me get to this stage.

I consider it a privilege to be here and be allowed to talk about acoustics. I have not done this for some time, as will become evident as the talk goes on. (Laughter)

Nevertheless, I have heard some great lectures the last two days and I hope that mine can keep up with the standards that have been set.

In the photograph that was taken yesterday I think you will have no trouble recognizing me -- Logan promised they would be printed before we left. You remember, "The man with the gray hair please move to the right one foot." I was thinking he could have said, "The man on the end of the second row please move, or the man in the dark shirt, could you move to the right." With Tom Gabrielson he said, "The man with the red hair please step forward six inches," and I thought that was very discrete. (Laughter)

I am the man with the gray hair. I was born in the year 1922 and I spent 10 years at Bell Labs, three years at Sandia, 10 years at Xerox. My research career, highlighted by the acoustic microscope, started when I had passed the age of 50. I think that there is still more to come; if I am invited back next year, maybe we can share some of that with you. But there is no question, the highlight was in the 1970s when we did the work that I will show you this morning.

I was asked how I got started in acoustics -- I was not trained in acoustics (I was trained in microwave electronics), but I was able to get some students and some very good people who did know acoustics. Ed Carome sent me Jim Havless, Izzie Rudnick sent me Joe Heiserman, and I benefited a great deal.

I started in this field when I was sent to Sandia in Albuquerque. Sandia was the engineering lab for nuclear weapons, atmospheric testing had been banned, they were going underground, and then they wanted to ban underground testing of nuclear weapons, and we were given the assignment of detecting underground explosions so that we could find out who was violating the ban.

There was a great deal of controversy in those days, because there was a school of thought that you could not detect underground explosions. Our assignment was to set up a system. We chose seismometers and I told the group, let's start a network of seismometers.

They went out to Lone Pine, California, which is east of here on the other side of the Sierras, and I did not look at the project for 4 weeks and then, when I checked up on it, I found it spent \$250,000 in 4 weeks. The number is just burned into my brain, because if you had extended it over 12 months, it was the entire budget that we had available.

Nevertheless, the seismic network was built and it is worldwide, it is still operating, and, as you know, it first missed the first real test of India a few weeks ago. It did not miss Pakistan, but the Indian explosion escaped detection. In my view -- and this is my view, I do not have any

data -- they did not explode as many bombs as they thought (not bombs, but tests) and the yield was not as high as they published, because the seismic network really works quite well.

The next thing, when I arrived at Stanford, I had to find a field of research and optics was a natural choice, lasers were very powerful, just coming into being -- they had Shallow there, who invented the blasted thing, they had Sigman, who was writing a book, and I did not feel I could make an impact against that competition.

Silicon was coming in, they had Lynville and that was sort of cornered, but no one was working on acoustics, so we started out with the interaction of acoustics and optics, acousto-optic modulators, and we did very well. You may recall that the laser printer uses acousto-optic modulators to make xerographic machines into printers. Both Xerox and IBM at that time went into that work.

We were working on the acousto-optic modulators for some time and the reason you can modulate light with acoustic waves at microwave frequencies is the wavelengths are more or less similar, but it took us a long time to realize that with a wavelength similar to the optical wavelengths you could make a microscope with a resolution equal, so we invented the acoustic microscope and decided to use it for imaging.

Allaroff is the director of the Physics Institute in Saint Petersburg and I told him that we had this marvelous invention, the acoustic microscope, and he said, "That's very good but you didn't invent it." I asked who did. "The Russians did." In those days the Russians invented everything. (Laughter)

I thought okay, but he was right. It started in the early 1940s with Sokoloff in Saint Petersburg. It is a very famous family there and he pointed out that sound in water at 3 GHz, 3000 MHz, had optical wavelengths and it should be used for imaging but he had no way of generating 3 GHz sound in the early 1940s -- we could hardly generate electromagnetic waves. The clystron had just been invented, the traveling wave tube did not exist, but he was a pioneer far, far ahead of the technology and we should not forget his early contributions.

Maybe we were the first in the 1970s to do the acoustic microscope and with our background on acoustic-optic interactions how would you set about to make a microscope? The only thing we knew was the optical microscope. We did not know much about the electron microscope and we were looking at waves rather than particles, so we copied to the best of our knowledge the optical microscope, which was a big mistake, I assure you.

[Slide 2]

It was a remarkable device. Just look at the top and you will see an optical beam coming down. You will see lithium niobate, which has an optical beam here, and you will have something here called an acoustic transducer that is sending microwave sound into the cell, liquid cell, that contains the object.

Zinc oxide with the transducer, 860 MHz was the frequency, and then we had matching layers and the object was in a liquid here, sapphire, YAG, niobate, mercury. You cannot imagine how long it took us to build this device.

[Slide 3]

How do you line it up, how do you get the image? We did manage and this is our first image of a grid, which we thought was absolutely marvelous, 50 μ across -- not exactly a microscope, because you can see these things with your eye. Nevertheless it was a start.

The thing that shook us was that we were sitting in the lab one day and Walter Bond, whom we will come back to later, looked at this and said, "Five surfaces to line up, all misaligned, all opaque, no way of seeing through them." I was sitting on a stool and I can still feel that stool that I was sitting on, Walter Bond, the sage, talking about lining up these opaque lenses, and I decided this is not the way to do it.

Many other things -- the wavelength of sound is in solids and solids have a high velocity, a large wavelength for a given frequency. Furthermore, the numerical aperture is not optimum, because these are plane waves, and if you want to get to optimum resolution, you want converging waves. There were many, many reasons why this was the worst possibility.

The other thing I remember is that we went to American Optical and had them run these lenses through their computer programs for checking aberrations and they were perfect, there were no aberrations, which we thought, well, okay, we did a good job.

[Slide 4]

In order to move away from this, we went to another system with Jim Cunningham, single-beam imaging, in which we had a liquid cell, plane waves of sound going out --

[Slide 5]

-- and here is a better view of it. Again, our transducer is plane wave. Then we had a membrane and on top of the membrane were latex spheres and the radiation pressure would condense the latex spheres into the pattern of sound.

[Slide 6]

Here is our first use of radiation pressure to condense 1- μ polystyrene spheres and reproduce the image that was formed by the sound passing through the object.

[Slide 7]

Here are grids, optical-acoustic. Here is a resolution chart, Air Force resolution chart, with 3250 lines per millimeter. We thought it was just marvelous stuff.

[Slide 8]

Here is lung tissue, optic at the top, and acoustic. Again, the resolution is not what you would like. It is contact printing, there is no magnification and there are a lot of defects.

In the meantime, Bert Ald at Stanford and Larry Kessler and Adrian Corpell at Zenith in Dr. Adler's lab were working on sound in liquids and they would image the surface of the liquids with scanning laser beams, a scanning laser acoustic microscope, SLAM, so we had competition, and this was not getting us anywhere.

We changed. The day we changed I was driving back on El Camino from Green's lab at SRI in which he showed me a big water tank he was sending sound through at low frequencies and imaging printed circuit boards.

In Palo Alto, which is the name of the town, they have a Palo Alto, which is a tall tree. I was driving and looking at this tree for some reason -- I had never seen it before and have hardly ever seen it since, but there is the tree in my mental vision -- and I thought about Green's images of the printed circuit board and I said there are no aberrations, there are no spherical aberrations in those images and I wondered why not.

By the time I got back to the lab I had figured it out, because if you take a spherical surface, go from quartz or sapphire into water, the ratio of sound velocities is almost 10:1. That means that the acoustic ray going from the solid into the liquid has to leave the surface at normal

incidence within a very close approximation. The focal length of a spherical lens in sound is equal to the radius of the sphere. There are no spherical aberrations, it is a perfect lens.

[Slide 9]

With that realization we changed course and we called upon Walter Bond, who had joined our lab -- Walter Bond is famous because of this device, which comes from the cover of Simon See's book on transistors; this is the first transistor ever made at Bell Labs, and it was made by Walter Bond in the 1948-1949 time period. Walter Bond was a mechanical genius, he made everything possible that followed. He made the lenses, showed us how to do it, and what we did beyond that is due to his skills and the skills of the students.

Ross Lemens, whom you have heard of at various times, went on to be a division leader at Los Alamos, was the first student who did the microscope in transmission at around a gigahertz using spherical lenses in sapphire-zinc-oxide transducers.

[Slide 10]

Some of the images that came out are like this. This is a tissue section of a cancerous tissue. We spent a lot of time trying to interest the medical profession in this. We thought that the contrast for the cancerous versus normal tissue was larger in sound than it was in optics and, indeed, it was, but with staining you could bring out many of these features in the optical image, so why bother with new technology?

[Slide 11]

Nevertheless, here are cells, living cells, in Lemen's transmission-type microscope.

[Slide 12]

Here is another living cell. We thought these were just fabulous, because we can see contrast of these fibrils and fibers going through and the nucleolus, and so on, in these cells was much larger than in the optical business -- and we will come back to this -- particularly in living cells, where stain was not available.

[Slide 13]

The next step was to go into the reflecting microscope and this will allow me to talk about some of the features. The transducer is zinc oxide, sends a plane wave of sound through the sapphire and you have a lens here that is spherical, has a matching layer, and the rays coming down vertically will leave this surface at normal incidence and come to a focus at the center of curvature.

There are no aberrations and this diameter is determined by the diffraction of the wavelength of sound. If you use water, you will find that the sound velocity is 1500 m/sec and the wavelength can be calculated from that. We made a lot of these lenses and did a lot of imaging with several very, very good students. Abdullah Attallar was the first one -- he has gone on to be provost in a university in Ankara, Turkey, and he still returns in the summer to teach us about our present project.

[Slide 14]

So this is what you can do. This is the sound image of a quarter. This is the front surface -

[Slide 15]

-- and this is the back surface; you are looking through the coin now. Bob Gilmore of G.E. -- I think this is a Swedish coin that attracted him for some reason.

Reflection -- acoustic microscopy in the range of 1 GHz was our first project and some of our first publications.

[Slide 16]

Then we started looking at materials. Why bother? What can you gain with these acoustic microscopes? This is a cobalt titanium alloy and you see, on the left, an optical image that is rather faded and structureless. On the right you see the different phases in this alloy, 59% cobalt, 41% titanium.

This is just fabulous, because now we are beginning to compete with the optical microscope -- not in resolution, but we are beginning to see something that you do not see optically.

[Slide 17]

Some brass, simple brass, highly polished, and you can see the grains and stripes in the acoustic image at the bottom that you do not see in the optical image at the top. These are about 50 x 90 u. Our scanners were very limited in those days. We are scanning this image, this spot over the sample to make the image.

[return to Slide 13]

You realize, if we go back, that the defect of this system is that there is no field of view; you get a response at only this one point. To get the image you have to move the sample under it, the scanning (it is mechanically scanned). This is called a confocal microscope. Later on, Rudy Coffner came into our lab and said we should do this in optics and he invented the confocal optical microscope, which is extensively used. We thought we invented it, but it turns out there was a fellow in Czechoslovakia who was far ahead of us, again.

It is very hard to make inventions for the first time. There is always somewhere -- somewhere -- 10 years earlier that beat you to it and you just have to learn to accept that sort of thing.

We are mechanically scanning now, looking at things you cannot see in the optical image.

[Slide 18]

Here is an epoxy optical at the top, acoustic at the bottom. There are glass fibers embedded in this for strengthening this material. This is becoming very exciting now. Students realize maybe they have a thesis and I am realizing maybe we have an instrument. Resolution is still modest.

[Slide 19]

This one has always attracted me. At the top you see an optical image of a surface in which we put an indent. You can see cracks going vertically up there in the optical image. The field size is 325 x 325. Acoustically you see these strain fields around the image when you look at the reflected sound. The elastic properties are influenced by this strain field. This is an aluminum sample with an oxide layer on it.

[Slide 20]

At this point we had various sources of funding. The National Bureau of Standards gave us money and tied us in with Hughes Aircraft, Wageland, and others there, and they started building microscopes. Again, I remind you, plane-wave transducers, zinc oxide, usually spherical lens. Sapphire is a high-velocity material. Liquid (water) is low velocity and low attenuation. And

look at the steepness of the angle of the rays coming in at the side; this is something that is quite unique.

[Slide 21]

If we look at this ray out here, it comes in at such an angle that it will excite the Rayleigh wave, or the interface wave, that travels along the surface, decaying exponentially into the solid, into the sample. At the same time it is called a leaky wave, because it is re-radiating at that same angle into the liquid, and that angle we denote by θ_c . When it reaches this point it will re-radiate and be focused and appear as an output on the transducer.

You can see 2 paths now that will excite the transducer. The central one will come down, be reflected, and go right back. These ones on the periphery will come down at the critical angle, excite surface waves, re-radiate, be picked up. The phase delay in this one, with the longer path, differs from the central, and they interfere with each other.

The image you get depends on the distance between the lens and the sample. That distance we call Z .

[Slide 22]

Here it is, again, denoting this critical angle at which you excite the surface wave, which is re-radiating continuously into the liquid, and that can be picked up by this lens. You are not on the focal plane but you are moving off of it and this distance, which we denote as Z , we call this response the $V(C)$ curve for lack of a better name.

[Slide 23]

Here is the response of the transducer as a function of the distance from the focal plane $Z=0$. It is a maximum when you are at the focal plane and then it has these nulls -- like this -- and you can see that if you are at a narrow beam, $\theta=30^\circ$ (that is the outside limit of the beam), it does not excite the interface waves, so you get no interference, but if you open up the angle to 42° , then you get these interferences. The initial work was done by Atallar, Rick Gramersing, and Ralph Waygon at Hughes.

[Slide 24]

Here we thought we had a way of looking at information that you could not see optically. This is inconel, optical in the upper left, optical etched to look at the grain boundaries, but the grain boundaries appear completely different when we move off the focal plane -- this is the $V(Z)$ curve, another invention. But not so.

Some guy at Brooklyn Poly -- and you will forgive me, I cannot reach his name right now - did it before us, but he was very, very helpful and analyzed a lot of our work. He did it for electromagnetic waves.

Nonetheless, this has proved to be very useful. You will note that our frequency is increasing slightly up to 2.7 GHz.

[Slide 25]

Inconel, again, the different contrasts you get at different positions of the focus -- the focal plane is on the right here. You see the scale of 20 μ , not particularly good, but a lot of information regarding elastic properties. You understand that the reason this appears with different contrasts is that the interface wave velocity is different for the different elastic materials.

[Slide 26]

Here is a recording head. At the top it is rather smooth. These are the magnetic grains and we thought this was incredible and we would certainly start selling instruments at this point. As a matter of fact, we did. We did it in 2 ways. Olympus came and looked at this technology and they built commercial instruments in Japan and Leitz in Germany came and they started building commercial instruments in Germany -- great friendships, great times during those heady days.

This is all in water, it is all in room temperature.

[Slide 27]

Larry Lamm took a piece of titanium and took various images at different Z positions, $Z=-1$, $Z=0.05$, $Z=0$ in green, put it in color, and then put it together in this way to present this image of titanium, which I think would be very difficult, indeed, to do optically.

[Slide 28 - unavailable at time of printing]

I am indebted to these students throughout the years, because they did, of course, all of the work, and some of them were very, very talented and went on to great careers.

[Slide 29]

Here are some colleagues -- we are in a conference in Japan in the late 1970s. There is Robert Gilmore, who was at General Electric, did nondestructive testing -- he said the great virtue of this instrument is not at the microwave frequencies, which we were pushing, but at low frequencies where you could do nondestructive testing and look inside materials -- Chen Syi, who is a professor at Irvine, kept going on acoustic optic modulators and, over on the far right, is Andrew Briggs at Oxford, who is pursuing this technology there.

[Slide 30]

This is the same conference -- you recognize the fellow (my hair was less gray then). I am in a hospital. I am in a critical ward of a Japanese hospital. They did not have barbecues at these conferences, they had banquets. At the banquets, you can imagine, there is stuff called liquor floating around -- it was a great party.

Walking out to the car that night, it was completely pitch black and I stepped off the curb, not knowing it was a curb, and fell, and my head hit something on the way down. I was so embarrassed, I jumped up to recover, so no one would notice, and was looking at the fellow right in front of me, reaching out to shake his hand, and then I just fell backward and faded, collapsed.

The next thing I knew I was on the ground, they would not let me get up, and the ambulance drove up and took me off to the hospital. By that time I just stood up too quickly, the sphincter shut off here, and I passed out. They made me stay in the hospital all night and the conference organizer was so concerned that he put someone in my room with me to make sure that I was okay. He spent the evening teaching me to speak Japanese. (Laughter)

The next morning the doctor came in to release me and I said, "Doctor, why wasn't there any pain associated with this incident?" He was a Japanese doctor trained in Ohio, very fluent in English, and he looked up at me and said, "It was probably the alcohol." (Laughter)

Anyway, that was Japan. This effort over the years was punctuated by conferences -- some of them were just great, 1984 was a great conference. Micro '84 was the big microscope conference in London, and then Varenno, the great ultrasonics physical acoustics conference organized by Professor Sette. That is when I heard about thermoacoustics from John Wheatley.

Lake Cuomo is in northern Italy, maybe 100 miles north of Milan, and we all came in from various directions. Wheatley flew into Milan with his bicycle in the back, in the cargo, and he got on his bike and pedaled 100 miles north. We can still see him coming up the highway, checking into this beautiful villa on Lake Cuomo. Others would take boats across the lake; Wheatley would get on his bike and ride across the lake.

We heard about chaos there, we heard about collapsing bubbles, and so on. It was a great year, but that is not what I came here to talk about. What did I come here to talk about?

[Slide 31]

Cells: We were continually trying to get this instrument into the biological world, into the medical world, because we thought the contrast for cells, living cells, was greater than optical. Living cells you could not stain and we were right, you could see organelles and contrast in cells that were not available in the optical images.

[Slide 32]

Here is an example of a fibroblast from the heart of a living chick. These are living cells, they are moving. The contrast -- this is the $V(Z)$, there are dramatic changes, the fibers going out, which gives them adhesion to the glass substrate, and it is rather vivid compared to the dull optical image on the lower right, which was phase contrast -- April 26, 1980, 1.7 GHz.

[Slide 33]

Here is another one. We find these contrasts as the cell would change. These are living cells, with the interference between the top of the cell and the bottom of the cell, and we thought that this would give us a lot of information on the elastic properties of cells.

It was not a very interesting question in biology and so no one has really taken it up for this purpose -- a few, but that is not the main use of the microscope today.

We will now talk about the resolution. We want to improve the resolution. The frequency is limited by the attenuation of sound in water. We were working with about 60 dB of round-trip attenuation. The lens radius was about 40 μ . We could get up to 1500 MHz, but it was hard to go beyond that unless we learned to make smaller lenses.

We decided that the cryogenic liquids had lower velocities. Helium has a velocity that is about a sixth that of water, 250 m/sec, and we were headed toward helium. We went first to argon and nitrogen.

[Slide 34]

This is Dan Rugar and here is 2 GHz, and the period of this grating is 4000 \AA , 400 nm, and I think B is the scanning electron microscope, which gives you a certain amount of detail, and C is the optical microscope, which has a faded look that we always enjoyed.

At the top you see the detail that is available in the acoustics; 4000 \AA is not very dramatic, but it was a start.

[Slide 35]

Here is the chick fibroblast, again -- Dan Rugar, optical, upper left, acoustic on the right, and the SEM on the lower right. We are now getting reflection from the surface. These are frozen, so there is a large difference in impedance between the liquid and the frozen cell, so you do not see the contrast that we see in living cells where the sound wave penetrates into the cell itself.

[Slide 36]

This is another view of the same thing, where you are getting a lot of the detail that you see in the SEM, more detail than in the optics, 400 nm in argon.

[Slide 37]

Here we are at helium, 840 MHz, 1.95 K. I drove to UCLA to talk to Izzy Rudnick before we launched this project. We would like to go into helium because it is the lowest velocity fluid that we can find and we would like to go down very low in temperature. He said, "You should hire Joe Heiserman."

Heiserman and I went to San Diego and bought a dilution refrigerator with a top loader, which was very innovative in those days, and it is still in use today by the people pursuing research in superconducting materials.

We spent a lot of time with helium, Rugar, Foster, Heiserman, and we started out with these images of just gradings of photo-resist and aluminum on glass or sapphire, 250 Å period, 100 Å is the line width, and you see, at 840 MHz, we are getting good contrast with these small details.

We are trying now to improve the resolution.

[Slide 38]

Here we are, an integrated circuit in argon -- Bob Bore and _____, were instrumental in some of this.

[Slide 39]

Here is the comparison at the top in an acoustic image of an aluminum line integrated circuit with a scanning electron microscope. The details here exceed those in the acoustic, but the contrast is very vivid up there. These are 2-u lines on integrated circuits. This is a long time ago; this would now be 0.2 u if you were to do it again today.

[Slide 40]

Here, again, in liquid helium, the contrast available. You see we could focus on the different surfaces, because the depth of focus was very sharp and then we would render these in color.

[Slide 41]

Here is myxobacteria. Dale Kaiser worked on these gadgets, these devices, animals, and this is more detail than he had ever seen before in the mix of bacteria, the contrast, the elevation. This is frozen in helium and you could not get this topography using the electron microscope.

[Slide 42]

We pursued this with smaller and smaller lenses; we got down to 11 u radius on the lens. We went up in frequency, 4, 5, 8 GHz -- 16 GHz was carried out by people at El Segundo Aerospace Laboratories. We got to a resolution in liquid helium, a dilution refrigerator, of 150 Å.

We were fascinated by the competition with the x-ray microscope, exquisitely fabricated -- they would get to 500 Å. Our best images in helium were a resolution of 150 Å.

[Slide 43]

To give you some idea of the problems we faced, let's look at the timing for a radius lens of 13 u. You have a $T=0$ when you start counting and out here, you go in and you get the first acoustic pulse off the lens and then there is another one here that is off the flat, which you try to get rid of -- this is the second one.

The information comes out in this point much smaller than the first acoustic pulse from the lens, this is the information from the sample. These are nanoseconds. Six nanoseconds separates this pulse from the one that is on the flat surrounding the lens. You try to reduce the size of that flat, but it is not very easy. This is work that is done by Hoddy Marilu.

[Slide 44]

This is the best we did. We were making zinc oxide, getting quite good at zinc oxide transducers, and we decided to make one at highest frequency possible. To do that we made a laminated structure. We would deposit a film of zinc oxide, tilt the substrate, deposit a layer of zinc oxide at an angle.

The Z axis for active zinc oxide is normal to the substrate. It is a polycrystalline but the Z axis always lines up, so as far as the Z direction goes, it is a single crystal. Then we will tilt it. There was no coupling, so we made an alternating structure where the period matched the wavelength.

We were able to get transducers that were at 96 GHz and we got them with maybe 15 dB of loss and at that frequency in helium the sound was 30 Å. If we had pursued it, we would have gotten a resolution that was under 30 Å, and I am sure that is possible.

The problem is that you are down at 200 mK, you are in a dilution refrigerator, which is a very tiny box sitting in a helium doer [phonetic]. It is expensive, it is difficult, and you would have managed that, but the problem is that the frozen tissue, the frozen sample, reflects the sound and you see only surface topography. That is all you can look at because of the high impedance contrast and there are better ways of looking at topography using the scanning probes.

The helium microscope was pursued for a certain extent but it never really made it in the commercial world; very expensive, difficult to make, and limited in its topography.

I think that we will take a break at this point and come back.

DR. QUATE: That more or less tells you what we did in our lab. In the next lecture I would like to tell you something about what is going on by others. Let me just pick up a few points that I missed on the first one.

Someone asked how do we scan the object. The scanning devices were driven by loudspeakers. The cones of loudspeakers gave us the motion we wanted using the electromagnetic coil and we would mount the sample mount on a flexural device that was 2 flat plates that would bend in unison to give us the parallel flat motion that we needed to keep the image in focus over the scan.

This is, again, Walter Bond, who was always using flexural devices to maintain the accuracy that we needed. Someone asked how do you grind the lenses. Again, Walter Bond taught us how to do that. He literally took a very sharp tungsten wire and ground the end off and then he would spin it at an angle and grind a hole in the sapphire just as though you were rotating the point of a pencil and grinding it down. They were very, very good lenses as far as sphericity and so on.

Others have been able to etch these lenses in silicon. If you use an isotropic etch, you coat first the silicon with silicon dioxide, open a little hole, and then introduce an isotropic etch and you will etch a spherical cavity underneath the oxide. This was a technique introduced by Kendall Weiss at the University of Michigan.

More recently, Yamanaka and others in Japan have used frenel [phonetic] zone plates as lenses that operate on a flat surface that have certain advantages.

The work was done simultaneously and carried on after that by other groups, Kessler and Adrian Corpell at Zenith. We have mentioned the SLAM scanning laser acoustic microscope. Kessler set up a company in Chicago and I think that is still ongoing.

In Sendai, Shubashi and Kushibiki did the line-focus microscope using a cylindrical lens that has many advantages over the point focus, and other groups, such as Briggs at Oxford we mentioned, and people in Switzerland, Italy, and Germany. The helium microscope, the high-frequency microscopes, as spectacular as we thought they were, never attracted commercial interest except in selected research labs around the world.

I would like to start this lecture by telling you what did attract commercial interest and then go on to a few topics that might point the way to the future. While we were working in the 1970s on microwaves and pushing to higher and higher frequencies and smaller and smaller lenses -- Marilu went to 5 GHz in water at room temperature and the resolution was around 2000 Å, which is very difficult to get optically, and then we went on up to 8, 16, 96 GHz.

In the background was the director of the Gingson Laboratory, Marvin Shirterall, and he said, "You guys ought to look at a way of doing something that will catch the interest of the commercial world," and he kept this theme ongoing. We said we do not have contacts with the commercial world but we know something about microwaves and, beside, we want to compete with the other forms of microscopes called the electron microscope and the x-ray microscope.

In a sense, we did, in the sense that I have mentioned, but what was it that caught the eye of the real world? I think that we can illustrate that in the next few slides.

[Slide 45]

This is a pattern of chrome on glass. It is chrome that is 1000 Å thick. At the top you see an optical image. The pattern is just for the sake of illustrating this effect. Down here you have the $V(Z)$ curves for two acoustic images, one that is -1 u, the other is -0.5. You see these dark edges around the edge of the chrome.

It is a problem of adhesion. The chrome is not sticking at this point and the adhesion of metal films to various surfaces is a very important problem. We published this and thought that was the end of it. We should have spent all of our time on this problem.

Others did. Ian Smith, the graduate from Dr. Eric Ashe's group at University College in London, did a lot of very good work on acoustic microscopy. Ian Smith went to work for Vacuum Generators --

[Slide 46]

Excuse me, I am ahead of myself. Let me just illustrate that the adhesion said we are looking underneath the aluminum, the interface between the chrome and the glass substrate. Here is another vivid example of a honeycomb structure that has a surface of poly. Optically you do not see much; acoustically you can see the honeycomb structure inside. You are looking inside the structure, we are beneath the surface, and that is where all of the interest is in the present day.

[Slide 47]

This is a picture that comes from Olympus -- we will get back to Ian Smith in a moment, I am ahead of myself. This is an image of paint. You are looking at the top surface of paint and

underneath you see these white areas that are rust spots, and there is a sort of bubble underneath that does not transmit the sound to the substrate and reflects it a great deal, so this is defective paint.

This was done at Olympus some time ago where there are commercial instruments. At the present time there is an institute of acoustics in Windsor, Canada, which is across the river from Detroit. Dr. Myev, a Russian who has pursued this acoustic microscopy for some years, has now come to Canada. Chrysler is funding his institute in order to look at the paint on the Chrysler automobiles. I think this is a major problem that they are doing, funding it, and Dr. Myev gets great credit for setting up this institute and pursuing this line of research.

This is sort of low-frequency stuff. Most of the work is done around 50 to 100 MHz, which was beneath our notice in the days when we were pursuing microwave high-resolution stuff.

[Slide 48]

Now we come to the work of Ian Smith at Vacuum Generators somewhere near London several years ago. They said the advantage of acoustic waves is not the resolution but the fact that it can penetrate solids, which does not come as startling news to this group. With the lens they wanted to look at something that was quite practical, such as integrated circuits that were packaged and bonded to a heat substrate (you understand the heat dissipation is quite a problem, so you bond it to something that will dissipate the heat).

[Slide 49]

Here you have, then, the bonded lid and some of the defects around the edges, courtesy of Ian Smith. This is a hermetic seal and you see defects that leak to the atmosphere and destroy the seal and contaminate the atmosphere. Field of view is 12 mm x 12 mm, using ultrasound.

[Slide 50]

Again, the heat-sink bonding and packaged transistor devices, and we are going to focus the beam on the bonding layer underneath the lid.

[Slide 51]

Here you see the voids and so on that you would see in no other way -- Ian Smith.

[Slide 52]

Field of view is 7.5 mm, copper heat sink for powered devices used in the I.C. industry.

Looking at adhesion and bond integrity underneath paint layers, looking inside the material is probably the greatest use of the microscopes in the commercial world today. Just within the last 6 months John Krogler at Hewlett Packard in Portland called me, trying to get his hands on a microscope to look at faults in his packaging of I.C.'s in the Portland plant of Hewlett Packard.

It is ongoing and some of the best work is done in Windsor by Dr. Myev. The Russians have set up to make these lenses and they are doing it rather inexpensively. If I wanted to get an acoustic lens today, I would go to Moscow, because they make them for less cost than in other parts of the world.

I would like to shift now and go on to things other than straight acoustic microscopy. You recall when we started we looked at something called acousto-optics, where we used acoustic beams to modify optical beams, acousto-optic modulators, a very common device.

Kumar Patel at Bell Labs decided that he could use photo-acoustics to measure the optical loss in low-loss gases. In this system he would heat the gas with a laser beam that was absorbed

in the gas, modulate it, and look at the sound that was created by the absorbed heat by thermal expansion. We call it photo-acoustics.

[Slide 53]

In our version we did it in this way. We would take an optical beam, focus it onto the sample, which we call substrate, and then image the surface with the scanning acoustic microscope. This device on the right is acting as a receiver, so it is not the reflection mode that we talked about before, but it is a transmission mode where the input is an optical beam.

Then we scan across the sample and pick up the thermal response that is converted to sound and pick up the sound wave. Clate Williams did this work in our lab. Many others are working on this system. In France they did beautiful work, at Wayne State, and others.

[Slide 54]

This is the comparison of the optical image of a pattern on glass; this is just a gold pattern on glass. The acoustic image is up on the upper right and the photo-acoustic is on the upper left. You can see certain defects and so on in the photo-acoustic, another mode of imaging that presumably gives you more information than either of the other two.

The person who has done some of the best work, the most successful work, the most commercial work, in this field is at Brown University, Roger Maris, who discovered the semiconductor industry. In the semiconductor industry the lines are shrinking and that has been the remarkable progress that accounts for most everything that has happened with computers, that is, the line width has gone down to where it is now a quarter of a micron.

To utilize these narrow line widths in the FETs you must use thinner and thinner oxides, so they are talking about gate oxides that are less than 100 Å in thickness. The conventional methods of measuring the thickness are very difficult to use as you get down in this thickness -- ellipsometry, and other ways.

[Slide 55]

What the group at Brown did was use photo-acoustics -- and you cannot read this because I copied it out of a semiconductor publication, but let me just go through it for you. At the top he has a film in blue and he is going to hit it with a laser pulse picoseconds in length, generate a thermal pulse at the surface that will convert to sound because of the thermal expansion.

The sound will propagate downward and reflect off of the interface between the blue and the orange. When it comes back, it will change the reflectivity of the surface and that can be picked up by the laser beam.

[Slide 56 -- unavailable at time of printing]

Here is a pulse. The time -- you see it -- is in picoseconds, 0, 100, 150, 200 psec. It echoes from a 5-layer stack. The stack is 180 Å of titanium nitride, 280 of titanium, 548 of aluminum copper, and 433 of titanium. This, then, is a very important area for the semiconductors. Semiconductors have an infinite amount of money to spend on instruments like this and this is probably the widest use of photo-acoustics today.

[Slide 57]

Here is another -- picoseconds. Now we do not have 100's, we have 5, 10, 15 acoustic vibrations in a bilayer consisting of 390 Å of titanium nitride deposited on PVD titanium. The blue line is the data, the red line is the theoretical fit.

This group out of Brown is putting other people out of business using the ellipsometer, because this is a better system when your films get down to 100 Å or so.

[Slide 58 – unavailable at time of printing]

Here is another illustration from this article by Maris in *Future FAB*. FAB is the fabricating facility that makes transistors, 200 Å-thick titanium nitride film on aluminum. It is a commercial instrument that is being more and more widely used.

[Slide 59]

Change subjects. We will go back to Mark Hamilton's favorite subject of nonlinear acoustics. In our system, we found -- Rugar found -- Dan Rugar was working on liquid argon, cryogenic microscopes, when he discovered this effect -- Rugar is now working at IBM Research Labs in Almaden.

He took a grading and it was gold on glass, something like that, and it had a period, a pitch, of 2000 Å. This is in liquid argon. The top image is an acoustic image of this surface of this grading at low power. You do not resolve the 2000-Å grading because the resolution at this wavelength is insufficient.

He turned up the power and suddenly the grading became visible. This is a nonlinear improvement in resolution. He went on to study the effect is real.

[Slide 60]

Here is a better illustration of this. We are starting, then, nonlinear resolution of a 0.29-μ grading liquid nitrogen. The wavelength is 0.43, so the wavelength is almost twice that of the grading period, 2 GHz in nitrogen. On the upper left we have a power of 10 dbm. On the upper right we get 15. You cannot resolve this grading. At 20 db it is beginning to appear and if you crank it up to 25 dbm, then you can clearly see this grading.

What we believe to be happening is that in the focal point of this highly converging beam the helium is highly nonlinear, you are generating second harmonic and the second harmonic content gives you the resolution and it converts back to the fundamental in the reflected signal to give you a response with the same transducer that you inserted.

Rugar has worked out this theory. It makes reasonable sense; it is not completely solid (not enough work has been done on it), but what I believe is that the nonlinearity in the helium is giving you second harmonic energy that is improving the resolution. You can get about a 40% improvement in resolution with this nonlinearity.

[Slide 61]

This leads into some work on other nonlinear systems that are of great interest to me in the present time. This is carried out by Walter Arnold at the Fraunhofer in Sauerbruchen. There is now a consortium of workers in Europe, in Sauerbruchen, Montpelier, Switzerland, London, and Berlin working on the subject that I am going to talk about in the next few slides.

I want to introduce another form of nonlinearity by looking at the transmission of sound through the bond between material 1 and 2, indicated by the black interface. We are going to have 2 frequencies impinge on this, one at 5 MHz and one at 7 MHz (this is low-frequency stuff we will get into later).

If you look at this, you will find that you get on the output a mixing that occurs at the different frequencies, some frequencies, harmonics, and so on. Arnold attributes this to poor adhesion at the bond, because if you have point contacts there is a nonlinearity in there and that

nonlinearity is exhibited by mixing the difference frequency and some frequencies. He looks at this and uses this as a tool for measuring the quality of the adhesion.

[Slide 62]

Here it is, again, with the 5 and the 7 and all of the harmonics and the difference frequency at 2 MHz. If you study this difference frequency, you gain a lot of information on the bond strength of this surface.

[Slide 63]

Greenleaf at Mayo Foundation has just published an article in *Science* using nonlinear acoustics where he uses 2 focused beams, F1 and F2. He looks at an artery and the acoustic emission. This is the artery. He is looking at the edges. You see the plaque that is on the edges. You see how the edges are pronounced in the upper 3 images, which are the difference frequencies between the 2 frequencies that he is impinging on this artery.

If you look at the C scan, look at the standard imaging where you look at the reflected energy at the same frequency that is impinging you get this confused mess. He attributes this, then, to the nonlinearity of this media in reflecting sound from the artery. This is a fixed tissue; he does not have images *in vivo*, but I think the nonlinear imaging of this sort of medical device, medical objects, will be more and more used due to the work of Jim Greenleaf at the Mayo in Rochester, Minnesota.

[Slide 64]

My favorite topic these days is based on this device, which is called a cantilever. You see a beam, a diving-board-like beam, not a beam of light, physical beam, coming out here, supported on the left with a sharp tip on the end. Some people call this the force microscope, some people call it the scanning probe.

Its virtue is that you can measure the deflection of this beam with extreme accuracy. You reflect a light beam off of it and with a suitable photodetector you can measure motion in that beam to a fraction of angstroms. Then you scan it over the surface to take the image and if you are looking at topography it has a resolution that can be as small as a few angstroms. Mostly it is used in commercial products where the resolution is a few nanometers.

[Slide 65]

We are interested in it because of the nonlinear behavior of this instrument, the cantilever and the tip when coupled with sound. I think the first to do this was Seonuck in Warsaw, who made something called the pin microscope that resembled this tip and he vibrated it and did studies on that. I do not have an illustration, but he is clearly the pioneer.

This comes from Yamanaka and Kolosov in Japan. Kolosov is now pursuing this work in Briggs' laboratory in Oxford. They have the force microscope, the cantilever supported on that with the tip, and they are going to vibrate the sample with ultrasound. If you had a standard force microscope and scanned the tip over the surface you would simply record the topography of that sample. That is the basic operational principle of the force microscope.

In their work they modulate the sample with an acoustic signal, 1 to 10 MHz, and then look at the vibration of the surface and measure it with the tip. When the sample is hard, the surface vibration corresponds to the vibration of the amplitude of the transducer. When the sample is soft, the vibration of the surface is reduced below that of the transducer and so you get a soft

signal, low, and then when you get to a rigid place, you get a large signal and then it will go soft, again.

This gives you a very good measure of the elasticity but the resolution is no longer determined by the wavelength of sound. These are low-frequency signals, 10 MHz. We would never, never have thought of using 10 MHz in the days of our glory in the 1970s when we were doing imaging.

The resolution is determined by the sharpness of this tip and that can indeed be very sharp.

This is the ultrasonic force microscope of Yamanaka and Kolosov in Japan. Yamanaka is still pursuing it in Japan and Kolosov is in London.

[Slide 66]

Here is the version that is used by Walter Arnold in Sauerbruchen and the consortium consisting of these other groups around Europe: A transducer, a delay line, and a sample, and the scanning probe going across the sample to pick up the vibration. The transducer is vibrating the sample with a certain amplitude and the tip is monitoring the amplitude that reaches the surface.

If it is hard, the amplitude here will equal the amplitude here. If it is soft, the amplitude will be reduced below that of the transducer. There are various other modes. You can get to higher frequencies by looking at the flexural modes but the principle is still quite the same of monitoring the surface amplitude over the surface when driven with a transducer with a constant uniform amplitude over the sample.

MR. WAKELAND: When you say that when it is soft the amplitude is not the same as the amplitude of the transducer, are you talking about just the amplitude of the tip? Is it actually gouging down into the soft surface or does the surface actually have less amplitude?

DR. QUATE: Let's assume that the piezoelectric transducer is moving up and down with an amplitude of 1 Å and if this sample were rigid, this surface up here would move up and down with 1 Å of motion; it would replicate what the transducer is generating.

If this were Jello, then the amplitude here would not be equal to the amplitude here. It would be much less, it would be reduced. They could pick up, then, the fact that this is Jello and this is steel by looking at the reduced amplitude on the surface.

DR. MIGLIORI: Where is the 10 kHz appearing? Is that the resonance of the cantilever?

DR. QUATE: That is just modulating the sound. We are looking now at the 10 MHz, which is the particle amplitude. Ultrasonic force microscope.

You see, what pains me in giving this lecture is that we had something to do with the force microscope in the early 1980s but we missed combining the force with the acoustic microscope. That was done by others. It should have been done by us, so you see a certain regret as I go through this. (Laughter)

[Slide 67]

Here are other modes in the cantilever. We make these cantilevers by the hundreds but we missed this.

[Inaudible question]

DR. QUATE: The tip will follow the amplitude of the surface exactly. If the surface moves 1 Å, the cantilever and tip will move 1 Å. Maybe you are thinking of this as a horn or something to amplify?

There is no amplification. The surface of the cantilever moves identically to the tip and the tip will follow the contour of the surface.

DR. HARGROVE: The point is that the tip is smaller than the optical spot that you can attain and that is how you are getting your super resolution.

DR. QUATE: Thank you, yes. The advantage of the system, the importance of the system, is that you can measure this with optical techniques to about 0.001 \AA and that is why it is commercially important.

[Slide 68]

Here are some images, the only ones I could put my hands on. These are standard force microscopes and here is the ultrasonic image. There is more detail. They are different features. They are not completely understood, but I just want to suggest that this might be the wave of the future.

[Slide 69]

If I were going back into this business, this is the experiment I would do. This is bound to have some interesting consequences in the future but I have not seen any images yet. You have the force microscope with the tip and you have a sample that is being vibrated and you vibrate both the tip by moving it up and down in resonance and the sample with another frequency, so you have F1 vibrating, the resonance of the cantilever, F2 vibrating the sample, and you look at the difference signal. It is a nonlinear signal that is bound to be important.

The reason is the force between the tip and the sample is a nonlinear function of the spacing. I would recommend this to you for those of you who are looking for new research projects. This comes, again, from Walter Arnold, Sauerbruchen, Fraunhofer Institute.

[Slide 70]

I am going to finish, with your permission, with this. This you recognize as an acoustic microscope. This comes from Professor Kim. He did this work at the University of Hawaii (he is now at Berkeley in mechanical engineering). He calls it a liquid ejector.

If you go back to the acoustic microscope and piezoelectric transducer, zinc oxide, it should be PZT, spherical lens surrounded by liquid, you can focus the sound on the surface. Now, rather than put a sample there, you let the surface be free and you will create a mound of radiation pressure of sufficient intensity that will push the liquid up -- it depends on the surface tension, it depends on the diameter of the acoustic beam -- and you will form liquid droplets in a stream emerging from the top.

You have seen this many times in commercials, drops being ejected from liquid surfaces. If you drop a pebble into a pool you can often eject droplets.

This is something that is going to be important, in my view, and that is why I take time to finish the lecture on it.

Here you see a spherical surface and it takes time to grind it or, if you etch it in silicon, it is not always possible to get it the right radius.

Kim has substituted the Frenel zone plate here for the spherical surface, which works very well and it is easier to control than the other spherical surfaces -- again, ejecting drops from liquid. Why would you want to do this? The answer is there are several reasons you would like to have small picoliter-sized drops, 10μ in diameter, ejected with exquisite control from surfaces.

The molecular biologists would like to eject DNA bases. The people in the semiconductor industry would like to eject photoresist. At the present time, when you spin on photoresist, 98% goes down the drain; it is an environmental hazard and you would like to be able to reduce that.

Some people put ink in here and use it for printers, although the main printer is a bubble jet where you heat and create a bubble and that ejects the droplet, but there are many applications where sound has an advantage in ejecting small drops from liquids.

In this case, if you are going to coat a 6-inch, 8-inch, 12-inch silicon wafer with resist, it would take you a very long time with a single droplet ejector, but you could build arrays and this structure is amenable to building several acoustic ejectors in parallel.

Professor Myev at Windsor has 8 or 9 acoustic lenses in parallel for looking at the paint of Chrysler automobiles. I could not get any material on that, because it is not yet released, but they do exist and he will be going further in this direction.

My recommendation is that arrays of acoustic lenses driven by piezoelectric transducers, hopefully PZT, will come into more and more use and I would like, Mr. Chairman, to end the lecture on that point and recommend it to those of you who are tempted to work in this direction.

Thank you very much.

DR. GARRETT: At the point that you were oscillating cantilevers and the force, as it changed, caused that nonlinearity in the restoring force and the frequency would shift, is that --

DR. QUATE: If you take a cantilever, any cantilever, and subject it to a force, a vibrating force, it will vibrate as the force. If you subject it to a gradient of force, that will change the spring constant and change the resonant frequency, so the resonant frequency is a measure of the weight of the force acting on this cantilever.

In magnetic force microscopes, where they put a magnetic tip on and use it to monitor a magnetic field, it is the gradient of magnetic field, and they do that by monitoring the change in frequency. It is a very simple equation if you write down the motion in terms of a gradient, so you are correct.

DR. GARRETT: But it is for magnetic, not atomic force?

DR. QUATE: It is used for both. It is called noncontact, in which you can monitor the frequency and scan the topography with constant spacing, where the tip is spaced above the sample. It is a beautiful system, and you can keep the spacing constant because the tip will stay in a region of constant gradient -- noncontact atomic force microscopy.

MS. ZOU: Even the force between the tip and the sample would be very difficult to analyze. The force is difficult to analyze.

DR. QUATE: It is difficult, but it is not too difficult.

MS. ZOU: Especially in the Yamanaka. He also claims that when you vibrate the beam, actually the tip contacts when ultrasonic frequency is much higher than the modulating frequency, so do you think it is really possible that we can pull out the elasticity of the material.

DR. QUATE: Back up a minute. He says when you vibrate the tip, it does what?

MS. ZOU: The ultrasonic frequency is much higher than the modulating frequency. Actually, the tip is in contact with the sample.

DR. QUATE: Yes, it is in contact part of the time.

MS. ZOU: Yes, do you think it is really possible that we can pull out the elasticity property of the material just by --

DR. QUATE: Yes, I think it is a very fruitful field. There are complications in the theory, but that will give in to future work. I think it is going to expand, the research in this area.

MS. ZOU: How about the subsurface?

DR. QUATE: There has been work -- if you had a void embedded below the surface, then you would see that with the tip, because the sound amplitude at the surface would be affected by that buried object.

MS. ZOU: Would it depend on how much the indentation is?

DR. QUATE: It depends on everything, yes. (Laughter)

To get a precise answer, you would have to deal with a certain geometry, the depth, the wavelength, and so on.

DR. HOFER: Just to clarify this, I have not heard anything about these force-tip things in a while, but some of the early ones were -- many of them now are noncontacting? So there is an atomic interaction force that is what? van der Waals? And then some of the ones you are talking about are contacting? There is a distinction?

DR. QUATE: There are several regions. As you bring the tip toward the surface, van der Waals will attract the tip to the surface. As you get it closer and closer, it will repel, because the electrons in the tip cannot occupy the same spaces in the sample. There is a region where it is attractive and you can operate in that mode, called an attractive mode, or you can work it in the repulsive mode, so all of these modes exist.

You can put an electrostatic force, a voltage, between the tip and the sample and you get an electrostatic force of attraction. In one case we are now working with field emission from the tip, Fowler and Nordheim field emission, and we can control the spacing by controlling that current because it is a very sensitive measure of the field and the spacing in terms of field.

MR. PETCULESCU: How fast would it take to scan a whole cell?

DR. QUATE: A biological cell?

MR. PETCULESCU: Yes.

DR. QUATE: It takes too long. One hundred microns is a typical area and you take a fraction of a minute and --

MR. PETCULESCU: But could you see diffusive processes going on inside the cell at the same time as --

DR. QUATE: That would be difficult. With the force microscope you have a chance because you can scan these tips much faster than you can scan with the acoustic lenses, because they are physically smaller.

MS. PETCULESCU: What is the _____ made out of and how do they make it?

DR. QUATE: It is etched out of silicon, so you put an oxide dot on the surface of silicon and you use an etch that will etch the silicon and not the oxide, so you have a vertical column, and then you keep on etching and it will etch the material under the oxide to a sharp tip and it can be very, very sharp. These are all etched out of silicon.

MS. PETCULESCU: So with the acoustic force microscope they can get information from inside the material?

DR. QUATE: Inside the sample.

MS. PETCULESCU: Deep? Very deep inside or just very near to the surface?

DR. QUATE: Not very deep. It depends on the sound wave. How does the sound energy reflect when it passes through an object in the center? How does it appear on the surface? The tip will accurately map out the intensity of sound on the surface; that sound will be scattered by inhomogeneities inside and how deep depends on the wavelength and the nature of the scattering and so on.

Yamanaka, I think, has done the most work on this and his publications are useful.

POROUS MEDIA

James M. Sabatier
National Center for Physical Acoustics
University of Mississippi

DR. SABATIER: Anthony and I showed up at Mississippi together and there were 4 graduate students in the class. We already masters degrees, so we started taking advanced quantum mechanics and, of course, you have to size up your competition when you are in class and I very quickly realized that the other two guys I did not have to worry about, but Anthony was a significant challenge, so I spent hours and hours trying to work problems so that I could do as well as Anthony. At one point Hank Bass told me I was spending too much time working problems and I would never get my Ph.D. if I did this. (Laughter)

At that time I started going to see Anthony after the class and I would say, well, here is how I am going to work the first problem and he would tell me how he would work the second problem, so then we spent half as much time working those problems, so it did go a lot faster.

[go to Transparency 61]

This is the research entrance to NCPA -- I did this for Steve. Steve thinks that C has something to do with the Bible Belt when he visits us, but he has never been in the research entrance (he always goes into the entrance for the Jamie L. Whitten National Center for Physical Acoustics).

First of all, if you look in the background, you will see only 2 automobiles back there. One of them is Richard Raspet's pickup truck and the other one is my car -- you can tell how little money I make, because you can barely see it in front of Richard's Ford truck (this must be a Sunday or something).

This is the only lab with windows in it in the whole building. When this building was being designed, I was just finishing my postdoctoral there and was asked to stay on and I had the opportunity to tell them what kinds of things I wanted. I said, well, I want a place outside, some dirt, so we can do some measurements without having to go somewhere, so I got some windows and behind that forklift there is a hole in the wall to let you run cables and computer lines and things like that through there.

When it was being built, the building construction manager did not know what all these holes were for -- it was the only lab like this -- and he wanted to know what it was all about, so they told him it was for Sabatier and he was a Cajun, blah, blah, blah, so it got named "Cajun Heaven." (Laughter)

So this end of the building is referred to as Cajun Heaven. If you go through that door you will find a lot of things going on with porous materials and you will even find Felipe Gaitan working in the office next door to me. Felipe is working on some sonoluminescence still and he is starting to get back into thermoacoustics, and he is now going to look a little bit at mass transfer in porous materials as well, so I will stop with that introduction and go on.

[Transparency 1]

What I am going to talk about is the interaction of airborne sound with the ground. That is primarily what I have made my career doing and somehow I have been able to support myself by doing that. Porous materials is a very large field. You could go into architectural acoustics, the whole subgroup of the Acoustical Society, and a large number of people focus on sound-absorbing materials. What I have to say is very applicable to that and, in fact, that is where most of it comes from.

Then you can go to maybe the geophysics community and there is a very large contingent of people who work on sound propagation in rocks. They do not worry too much about pore properties of these rocks; they are mostly concerned about grain contacts in rocks. And then you can go to the underwater acoustics community and then you have all of the ocean bottom, which is a sediment -- I guess there are some deep ocean profiling people, but if you look just at the shallow first meter or so of the ocean, that looks a lot like the ground, except that it is water-filled, and then there is the ground.

I am going to talk about these 4 things, or at least I am going to try to. I literally mean that we could go take the kinds of measurements I am going to do down by volleyball net and when you get done this morning, you will have a different notion of the sand that we have been playing volleyball on. Or you could go out to the beach and take measurements on the beach out there, where maybe you notice the sand is a much different grain size and the acoustic properties of the beach sand are a lot different.

I have had the luxury of taking measurements in Panama City on a beach there near spring break. It is a very lovely time to visit Panama City, as many of you know, beautiful beaches.

I would like to talk about porous medium properties, the physical properties. One of the things I find a lot of students and colleagues of mine are always frustrated with is you talk about tortuosity and you talk about permeability and you talk about pore shape and they do not have a feel for these kinds of things and they find that kind of disheartening because of the intense number of parameters that you have to deal with.

I will go through those in some detail and I think when I am done you will feel comfortable and you will be happy with them, because I will show you that we can measure these. There are techniques for measuring them statically but we have developed very nice acoustic techniques for measuring tortuosity and for measuring flow resistivity of porous materials, and other people have done this in other communities as well.

Of course, sound propagation in porous materials is built upon sound propagation in a rigid cylindrical tube; that is how you start. It is just a rigid tube filled with a fluid or a gas in the case I am going to talk about today. Then, from there, you build up to bulk materials and talk about propagation in bulk tubes, and tubes are an excellent model for the ground, cylindrical tubes.

[Transparency 2]

When I first met Hank Bass, this is a slide that he showed me. This is a measure of the normalized real and imaginary parts of the surface impedance of the ground. The data go from 150 Hz -- actually, about 200 Hz -- out to 1500 Hz.

The way Bass measured this -- there are several ways you can measure the impedance of the ground. Of course, one of them is to take an impedance tube and stand an impedance tube up on its end and pound it into the ground and then make measurements where you look at the

standing wave profile and measure the maximum and the minimum and the distance between this maximum and minimum and back out impedance.

It does not work very well when you take B&K impedance and try to use it on the ground, for several reasons. The impedance is pretty high. When it gets above 10, B&K impedance tubes do not work that well because their reflection coefficient is pretty close to 1, so you do not get much difference, you do not get very reliable results.

What Bass did here was to set a loudspeaker up (this was outdoors in a big grassy field over about 100 yards) here and put a microphone in front of it, things are about a meter off the ground, go 100 m away and set up another microphone.

This is your source, and this is a reference microphone, and this is a measuring microphone. Then argue that the sound travels this way and it also goes that way. The sum of the sound at this microphone is due to spherical spreading of a ray going this path, $1/R e^{ikr}$ and then some reflection coefficient times 1 over another R, some larger R e^{ikr} times the reflection coefficient.

What he did was to say I can write the plane-wave reflection coefficient (and he did the spherical wave reflection coefficient) in terms of a complex impedance and he measured the magnitude of the phase right here and he twiddled the real and imaginary parts of the impedance until he matched the pressure, and that is how he got these data right here.

There are several things about these data that are interesting. First of all, if you said the density of the ground is 1000 MKS units and the speed of sound in the ground (and no one quite understood what the speed of sound was in the first centimeters of the soil at the time), you took a guess, you looked in a book, and you found out from geophysicists it is going to be 1000 or 2000 m/sec, so that gives you a ρc on the order of 10^6 of characteristic impedance.

The you take the characteristic impedance of air, which is 400, and you calculate that you are not going to get any energy into the ground or that the impedance is going to be infinite. The first thing is that the impedance is finite; it is only 10 times greater than that of the soil. That is the first thing that you learn about this.

It tends to look as if the real and imaginary parts are equal when you do this and you might have guessed initially that maybe it would be mostly real, mostly resistive. The other one is that the imaginary part of the data, if you extrapolate it out, starts to go to 0. In these data it is not so obvious, but if you run these data up to 4 or 5 kHz, you find that the imaginary part of the impedance gets very close to 0, so the ground does become pretty resistive and you do get quite a bit of absorption. The imaginary part is going to go to some value, non-0 value, typically on the order of 2 or 3 or 4.

Hank, being a good Ph.D. adviser, gave me those data and said go explain it. Go explain it in terms of sound coupling into the ground. I would like to know what is going on. That is how I got started doing what I was doing.

[Transparency 3]

I was particularly interested in getting to lower and lower frequencies. That meant bigger and bigger speakers. I also was interested in getting to normal incidence. If you get to normal incidence it is often easier to understand problems. You do not have to worry about, for example, shear waves if you are going to treat the ground as a solid, and you do not have this

cosine of the angle floating around. There are some advantages to getting to lower frequencies or getting to normal incidence.

Low frequencies meant we had to build bigger speakers and we built very large speakers over the years to look at the coupling of sound into the ground. To get to high angles we ended up using cranes to raise these things up off the ground, often very large cranes to get these things up as high as 150 feet for some of our measurements and eventually to large towers that you can find out at White Sands -- they have these kinds of towers and speakers erected up there.

What we did was design a microphone that we could push down into the ground. Everybody put geophones in the soil and once we decided to start thinking about the ground as a porous material, it would have been nice to have been able to get the other side of the porous material, but you cannot do that, so we designed microphones that we could push down into the soil.

In the earliest days when we started sticking microphones down into the ground, we used B&K microphones and they did not work very well. You drill a small hole and put your B&K microphone in the ground, you get some grit between the diaphragm and the grid and the microphone quits working, so you have to figure out ways to do that, so we started learning to put them in condoms and finger cots and things like that, which is another whole interesting story when you are a graduate student getting your degree in the Bible Belt and you had to go to buy your condoms at the drug store and then come back and give to the secretary and ask to be reimbursed. (Laughter)

You always have to write an explanation of what you are spending the money on. We eventually put these in tubes and later on I will show some of these.

We also took geophones -- since we put microphones in the ground, we also put geophones in the air, because we were concerned about how do geophones behave when they are surrounded with the gas (Tom alluded to this in his talk). A geophone is an accelerometer, it is a mass on a spring, but you get a velocity out of it and that is because the mass on the spring is a magnet and you wrap a coil of wire around the magnet and, when the mass moves in and out, you get an induced voltage and that voltage is proportional to the velocity of the case.

They are very good low-frequency instruments and they are very sensitive down at low frequencies. We were, hopefully, trying to get down as low as 10 Hz. They get quite large. If you get a geophone with a resonant frequency of a hertz, it tends to get quite big, on the order sometimes of a foot in size.

This is the measurement geometry, what I am trying to show. We can change the angle of incidence of the loudspeaker. We can push this probe down at various depths. We have a geophone that we can move up and down as well.

[Transparency 4]

Ignore the lines. What I want you to do is focus on the data. These are some data from Mississippi from what is called pit run sand, for some reason; it is a sand pit that you can find outside, the sand is unconsolidated, it has lots of clay in it, and we were able to make some measurements in this pile of sand at various depths.

Here is the probe microphone -- this is actually the attenuation -- and this is the probe microphone down at 5 cm. There are multiple measurements here. Then it moves down to 10 cm, 15 cm, and then up to 20 cm. If you look at those data at, say, 100 Hz, you have about 1 dB

of attenuation per centimeter. That does not seem like a lot, but if you change that for the fun of it to dB/m, you realize that the sound does not go very far in this material; it is highly damped. We see we have a lot of attenuation.

I was sitting in E&M class at the time that I was making that measurement and someone was discussing Kramer-Kronig relationships and I wondered if I could use those to figure out what the velocity was in porous material. It was probably one of the first times I learned to make some assumptions, because I had attenuation data over a rather narrow band and, of course, Kramer-Kronig is for a whole frequency, and I thought, well, I will just do it over this narrow band and extrapolate it out, and I got a number of 100 m/sec.

[Transparency 5]

I said, well, this should be measurable. These are the data that I measured, measuring the phase velocity using a probe microphone in some sand. The flow resistivity of this sand -- we used to use rayles/cm, but the reviewers do not like rayles, because they are currently not S.I. units, so it is actually 85 kN sec/m^4 (and I will go through flow resistivity in a couple of minutes).

There is a sound source sitting up here above the surface of the sand, broadcast the tone, put the microphone down at 2 distances, and just look at the phase between those 2, and then extrapolate the phase feed. That is what you get for frequencies from 50 Hz to about 400 Hz. There is a lot of dispersion and very, very low wave speed. For me, that was quite surprising when I first saw that.

The goal, then, was maybe to try to set out and understand that in terms of porous media models.

[Transparency 6]

This is what happens if we put a geophone down in the ground -- forget about the dashed line, that is a calculation and I will come back to those toward the end of the talk. The solid line is a measurement. Again, we broadcast a sweep of frequency with the geophone buried at some depth below the surface and we plot the velocity at the geophone divided by a microphone at the surface to take out fluctuations in the sound-pressure level of the loudspeaker. The pressure of the loudspeaker goes up and down; the geophone and the microphone are kind of co-located (the microphone is on the surface and the geophone can be buried below the surface). That is the magnitude of the coupling.

If you did a calculation of how much coupling you would expect to get into the soil just based on a plane-wave incidence on the surface of the ground from air into soil and you again did this naive calculation, you looked in the books and you checked to find out what the velocity was of compressional waves in the soil, and it was 1000 m/sec, this coupling would be 2 times larger than what you would expect. You could not find very good measurements of sound speeds in the first few inches of soil at the time.

Typically, when you do this measurement, you get this series of minima and maxima, which obviously have something to do with reflections, and it is because the ground is layered.

A couple of initial observations: If you change the angle of the sound source, these minima and maxima move around (kind of what you might expect from a layered system). If you do a time-of-flight measurement from the source to the receiver, you find that the sound comes from the speaker directly to the receiver. It does not go from the receiver down to the ground, for

example, and then travel at a high speed toward the receiver. The coupling of the sound into the ground appears to be local, it appears to be coupling in right above the geophone.

Those are just some data to set a background to give you a feel for the kinds of things that we are going to attempt to explain in the rest of the talk.

DR. HARGROVE: Did that geophone have a vertical component?

DR. SABATIER: That was a vertical-component geophone but we have certainly used radio components where we looked toward the loudspeaker as well. In earlier days we used triaxial geophones, because we did not know any better -- you could spend a lot of money, as you buy triaxial accelerometers, you can buy triaxial geophones, but you very quickly learned that they couple between each other and they are not very good in terms of making these kinds of measurements.

[Transparency 7]

The sample material consists of a frame or a matrix that is saturated with a fluid. ρ_s is the density of the solid component. The fluid density is ρ_f . We can define, then, the porosity of the system to be the volume of the fluid divided by the volume of the bulk material and, of course, the fluid plus the solid volumes are equal to the bulk material volume. We can define the density to be 1 minus the porosity times the solid density plus the porosity times the fluid density.

One of the assumptions that we make in modeling propagation through porous materials is that there are no isolated voids. There are no voids or pockets that are sealed off to the rest of the material. When we are going to talk about acoustic oscillation or acoustic flow in a porous material, we want to be able to have flow through all the pores.

[Transparency 8]

Flow resistivity: This is a very simple measurement, a trivial measurement to make. It is probably not done by very many people. You can look at new modern instruments for measuring it, but I am going to show you an old instrument that I prefer to use to measure flow resistivity. It is simply defined to be, if you have a sample in a tube, you set up a differential pressure across the links of the sample, and you would like to know the difference in pressure divided by the flow velocity.

What you really do in a measurement is to take the difference in pressure divided by the volume velocity for a unit cross-section area of the sample and that is the flow resistance. If you know the length of the sample, then just like resistivity or resistance in electricity you can divide, then, the flow resistance by L and end up with what we call the flow resistivity, and this has units of Newton seconds metered (-4).

An apparatus for making this measurement was developed by Leonard in the late 1940s, and it is published in JASA. Not long after that Isadore Rudnick came out with a modification to it, a fix to that apparatus (and I am sure all you Rudnick students used it and know very well what that was).

[Transparency 9]

What you do is you go into the attic of the physics building and you find an old beam balance. Some physicists never throw anything away, so you go to what I call the bones and you get the old 2-armed beam balance. I am trying to represent that by this beam right here. There is a base support that supports this thing and there used to be a pan over here and there used to a

little pan over here, and there is a little man that you can move back and forth here. This thing is calibrated in length, so you can have a little rider to adjust the mass (nobody uses these kinds of balances any more).

Here is the sample holder sitting right here. The sample holder attaches to a cylindrical fluid bath filled with kerosene, so this is a cylindrical channel, if you wish, that is filled with kerosene. Then a piston, if you wish, shaped in this way is set so that it can go inside the kerosene bath. This makes an air seal right here and then we can apply a force right here, so we can create a differential pressure between here and there, depending on whether we add weight to this side of the pan or add weight to this side of the pan, so we can actually pull the air through here.

We then hook a pointer to the center here -- in fact, the balances had a pointer, these old balances, it was already there, and there was generally a scale, so there is already a scale right here -- so you have a pointer coming down and you line this up so that when there is some high H you can move this pan up and down and keep this in the fluid.

If you pull this all the way down, this thing just usually rests on the bottom here -- it rests on the lip right here -- which, if you go all the way up, the pan just rests in the water here and you splatter kerosene everywhere (it is kind of messy).

You have to know this distance right here and how that corresponds to that distance right there. You can do this on this scale; you can make this very accurately. You can measure a broad range of flow resistivities. You can have very, very short times or very long times and you can easily modify this to expand this scale so you can measure small distances or large distances.

To get the pressure, then, the pressure on this side is just MG divided by that A_0 right there and the pressure on this side is atmospheric pressure. The sample has some length A . We measure the velocity of this thing coming up and that gives us the linear velocity. We know the area, that allows us to get the volume velocity, and you use a stopwatch; you just time how far it goes back and forth.

MR. POESE: Why do you use kerosene?

DR. SABATIER: You would like to have something that has a low viscosity, low friction on the walls, and that is the reason that you use kerosene. Water would probably work just as well -- well, I do not know how well water works.

MR. WAKELAND: I was kind of surprised in the definition of S that there is not anything to account for the properties of the gas. I mean, I was expecting some fancy nondimensional constant kind of thing to be in there that would just tell me about whether or not I was using helium or air in there or whatever. It is going to be different, isn't it?

DR. SABATIER: Yes, it depends on the fluid. Air flow resistivity, right.

DR. MIGLIORI: So is it time-dependent?

DR. SABATIER: No, it is not, as long as you -- we have to be careful. What you do is you first make sure that it is linear. That is, you make sure that the time and the weight, that this velocity is linear over the region that you are working, so you want to work in the linear regime of the flow, so it has to be streamlined flow, you do not want to create turbulent flow.

DR. GARRETT: The experiment to mimic the Reynolds number are the acoustics experiments?

DR. SABATIER: No, what I am going to do a little bit differently is to make the acoustics experiment look for the D.C. limit of this when I develop a model.

DR. GARRETT: So it is always low resonance in the flow?

DR. SABATIER: Yes.

When you set this thing up and you do not have the sample on it, this thing just sits here and oscillates back and forth; it essentially goes for a very long time. The reason is that as this mass starts to go in, it experiences an increased buoyant force and sort of pushes it back up and it starts to come out, starts to get heavier, so the thing just oscillates back and forth.

What Rudnick did was to say put this small mass on here and now this will counteract, then -- this comes over here -- will counteract the effect of the buoyant force and push it back and you can damp out these buoyant force oscillations very easily. Without that, it is very difficult, it takes a very long time, it is just like a pan balance without having magnetic damping. It is analogous to magnetic damping.

[Transparency 10]

Tortuosity is another measurement (which I have never measured, at least not electrically), which is an electrical measurement. What you do is take the porous sample, fill it with a conducting fluid and measure the resistivity with the conducting fluid and then measure the resistivity without the fluid, take the ratio of those 2 and the tortuosity is then defined as Q^2 as omega times the ratio of those 2 numbers. Apparently electrical engineers make this measurement a lot. We have had materials sent off to have the electrical tortuosity measured so we compare them with acoustical measurements and we get very similar results.

[Transparency 11]

There are 3 physical properties, then, of porous materials that one can measure in a D.C. sense. You can measure the tortuosity, you can measure the flow resistivity, and you can measure the porosity. I did not say how to measure porosity, but there is a whole host of ways.

The simplest way is to use a material in your experiments that you know the mineral content of and then you calculate the porosity, but there is a host of other measurements that one can use.

I like to think of the ground as soda straws, just a big box of soda straws that you buy at the store, and the only openings now are in the tubes themselves; the gaps or triangles between the cylinders is a solid so we have a straight tube model of the ground.

Maybe these have a diameter of $2R$ and we will assume an identical circular tube of radius R . Let there be N pores or radius R /unit cross-section area, so the porosity is just $N \times \pi R^2$, where N is the number of these pores.

Other kinds of models that you can use, and people have done this over the years, is to look at stackings of spheres and simple cubic hexagonal closest pack, et cetera, but it gets very difficult to try to model the pore space, and that is the problem with these kinds of models. It is almost impossible to come up with a way to model what the pore space looks like.

What we will learn as we go through this is that it is not that important, how you actually defined or came up with the shape of the pore space, what geometry you used. Random packing of spheres is a nice assumption to make about how sand or grains of soil pack into soil. If you have identical spheres, you know what the porosity is for those particular cases.

[Transparency 12]

The next 3 slides are only to give you a road map for where I am going to go theoretically. These are certainly approximations to Mark's talk yesterday, but I just write down the 3 linear equations of acoustics: state, continuity, and motion.

[Transparency 13]

Once one does that, one can then write down a wave equation. C^2 is defined to be the sound speed, the ratio of the bulk modulus divided by the density, and you can write down the general solution. The main thing here is that K is ω/C or ρ_0/β to the root $1/2 \times \omega$.

[Transparency 14]

If there is absorption in the material, then we usually write K as some real part plus an imaginary part, $K + \text{or } -i\alpha$, stick that into your plane-wave solution and that imaginary part $i^2\alpha$ comes out, so you get this attenuating term out in the front, the plane-wave back term -- this is the real part of the wave number.

We can rewrite the wave equation, which is why I did this now, so that the wave equation now is written as $1/C^2$, where C is a complex number, K is ω/C complex, where now the goal is, in a porous material, to write an effective density and an effective bulk modulus to include the effects of viscosity, viscous friction, and thermal conduction.

That is typically what has been done in the literature, it is what was done by Zwicker and Costin in the early 1940s in their book on sound-absorbing materials. Beale did this when he developed this scaling or developed -- Beale never referenced Zwicker and Costin, but Beale did the same thing when he looked at sound attenuation in porous rocks and considered the rigid tube to arrive at this.

What I am going to do now is to develop notions of what ρ and β are in tubes.

[Transparency 15]

We would like to get a simple understanding of what viscous friction is, so we envision 2 sheets of a fluid moving at velocity U_x and velocity of U_x plus the partial of U_x with respect to y dy separated by distance dy and, of course, Newton's third law says that these 2 forces are equal and opposite.

We would guess, then, that the size of that force would be proportional to the difference in velocity to the area of the plates and the separation distance, plug that into the equation there and you get U_x -- and I am missing a minus U_x right here -- times $dx dy 1/dy$.

The stress, then, force per unit area, is proportional to the difference in velocity -- you can, of course, can write this as a strain rate, $\Delta L/L \times 1/T$, kind of a physical intuition as to what this is on this side -- the stress, then, is the constant of this proportionality, the viscosity of the fluid times this gradient in velocity perpendicular to the surface.

Of course, if you do this for a volume, we have to remember that we have to include rotational effects of this volume, the shear stress, or the shearing of it, so this term actually has another term right here.

[Transparency 16]

If we let one of those plates oscillate back and forth at a frequency ω , then we can write down the viscous penetration depth, which is $1/E$ and the dropoff of the velocity from the wall, and for air at 10 kHz the viscous penetration depth is 20 μ .

The ground is made up often of grain sizes that range from 10 μ all the way up to several hundred microns. The pore size is typically on the order of a third of the grain size, so that we

see we can go from grain sizes, where the viscous penetration depth would be quite small compared to the pore size, to grain sizes where the viscous penetration depth will be quite large.

[Transparency 17]

Sound absorption in a single cylindrical tube with viscous drag and thermal conduction -- Zwicker and Crostin do this -- they initially assumed that the tube wall was rigid, Z axis is down the tube, R is the radius, and what they assumed when they initially developed this, was the pressure was independent of R and the frequency is below the first nonplanar mode that goes down the tube.

Since that time all of these assumptions have been verified by several papers, including, possibly, people in here. Tijdeman was the first person to look at the full-blown problem of sound propagation in tubes of all kinds of sizes and Pat Arnott did this, I think, in thermoacoustics, I think Greg Swift did this, Michael Stenson at NRC, so a whole host of people have repeated, essentially, the work that Tijdeman did for their own particular applications.

They showed that the assumption that the pressure is only a function of Z is valid; that is the main thing that they want to show. That is generally the starting point of all of Zwicker's and Costin's work.

The goal, then, is to build up a notion of how to incorporate the flow resistivity, σ ; the porosity, ω ; and the tortuosity, Q^2 , into bulk materials. That is what we want to do.

[Transparency 18]

You apply this to a cylindrical tube of radius R. We know, then, that we have plane waves, and we assume U_x is equal to U_y is 0 and the boundary condition at the wall tells us that U is equal to $U(z)$, U_z is 0, and r equal to R, and we like to look at Newton's second law, we like to look at the momentum equation. The momentum equation is $i \omega \rho_0 \times U_z$. That is $m a$ term. That is equal to minus the gradient of the pressure, the partial of the P with respect to Z plus A to the viscosity times this term right here, which is just the stress: It comes from the off diagonal component of the stress tensor for this dimension of motion in this cylindrical tube.

If you write this in cylindrical coordinates, you very quickly recognize that this is Bessel's differential equation, for which the solution is written down here, where the argument of the Bessel function LR -- where the argument is LR -- and L is $i \omega \rho_0$ to the $i \omega \rho_0 / 80$ to the $1/2$. It does not matter whether you use the positive or the negative root of L simply because of the symmetry of the Bessel function, about $\rho=0$. R is the radius of the tube.

[Transparency 19]

The next step is to integrate, then, this velocity to get the average velocity across the tube. We get the result, U_z is minus the expression that you see right there, where S now is defined to be the shear wave number, and its root, $\omega \rho_0^2 / \eta$, works out to be R times the square root of 2 over the viscous penetration depth, $\delta \eta$, and S is defined to be the shear wave number.

We can stick that back into the original equation and then write minus the grading of the pressures, $i \omega \rho_0$ times U_z , where ρ now here is referred to as the effective density, the complex density.

DR. MIGLIORI: You have a shear wave speed in here but you have been using a fluid?

DR. SABATIER: No, I did not mean to say the shear wave speed. It is the shear wave number. There is no shear wave here. S is the shear wave number and the reason it is called a shear wave number is because it is dimensionless.

MR. WAKELAND: Some understanding happened there, but I was wondering if you could --

DR. SABATIER: There are no shear waves. We do not really call these shear waves, all right?

MR. WAKELAND: But then Dr. Migliori said, well, yes. Now, why did you say yes? What does that have to do with shear waves? Why is it called it the shear wave number if there are no shear waves?

DR. SABATIER: It is the shear wave number because it is dimensionless and it has the shear viscosity in it.

DR. MIGLIORI: If you have viscous properties introduced into this, then the definition of a viscous penetration depth is the characteristic length over which a shear displacement damps out; therefore, it has a completely complex shear propagation that is omega, it is pure imaginary, and he can certainly define a shear wave number, but it is --

DR. SABATIER: Yes, and we will see that in a few minutes.

That is for a circular tube of radius R. You can do this for a slit and Biot did this. When Biot first did this, or maybe only when Biot did this, we started off with a tube that looks like this, so now what you want to do is to say, well, let's have a slit like this -- this extends off to infinity, plus and minus infinity in both directions -- and this now has a diameter of $2A$.

Biot's notion was that we need only 2 kinds of pores, circles and this kind of infinite circle, if you wish, and every other pore is going to fit between these 2. But unfortunately, what I think has happened is that there are a lot of theoretical people who said, oh, hell, let's look at square pores, let's look at triangular pores, mainly because they are workable. In my mind, that kind of confuses the issue in terms of understanding what is going on in porous materials, in bulk porous materials.

This is the result, then, of doing the exact same analogy for this slit, so we end up with this expression for this effective density for the slit.

S' is on the this slide -- I rewrite this one right here -- S' has A^2 for double the slit radius. We end up with S for the tube looking like this and we end up with S , the shear wave number for the slit, looking like that, so they are identical.

It would be intriguing -- what was done was to look at these 2 functions to see what they looked like.

MR. SMITH: Is conduction not important?

DR. SABATIER: It will be, but it was assumed that they could be separated and handled separately, and we will look at heat conduction in just a minute.

[Transparencies 20-21]

This is the ratio of this complex density or this effective density for a circular cross-section as a function of S . We see that S is what it is and we show the real and imaginary parts of this function. They are what they are, but if I look at the same thing for slits, I get very similar functions, so numerically they can be scaled.

This is what Beale says, that for this viscous drag effect I can scale between parallel slits and here, Pat Arnott did for square pores and Stenson did for triangular pores and people have done a host of other kinds of shapes in terms of doing this kind of stuff.

[Transparency 22]

I simply write down the results for the effective bulk modulus, which includes, as you said, thermal conductivity to the pore walls. This was done by a host of people; Attenborough did it, Zwicker and Costin did it, Craig does it in his Ph.D. thesis, Stinson did it in an article in 1991, and Tijdeman does the same thing.

For a circular tube, then, you just write down the thermodynamic equations and you can arrive at the bulk modulus in the same kind of analogy and these are the 2 expressions that you get for a circular tube and for a slit. Prandtl number takes care of the heat conduction for us; Prandtl number is $A (C_V)/K$, which works out to be about $3/4$ and for air it is nice, it allows us to make some approximations when we look at this. S and S' are defined as they were before, the shear-wave number for a circular tube and the shear-wave number for a slit.

[Transparencies 23-24]

If we look at what these 2 functions look like, that is, this effective bulk modulus for a slit and the effective bulk modulus for a cylinder, this is for a circular tube (you can see only half of it), and if we do the same thing for a slit, we get something like this.

[Transparencies 25-26]

Before I do the scaling, I would like to look at high- and low-frequency approximations for these functions. For small and large S we use the asymptotic approximations for the ratio of the Bessel functions for the effective bulk modulus (Prandtl numbers are approximately equal to 1, so you can use the same results).

For a large S , for a cylindrical tube, it tells us that the viscosity over $\omega \rho_0$ is going to be small compared to R^2 . That indicates to us that the shear-wave number, $\delta \eta$, is small compared to the tube size, so the effect of the viscous boundary layer goes away very near the pore wall, very near the wall of the tube, so the velocity profile in a tube is flat. That is our understanding of that.

If we make this approximation for Newton's second law for the momentum equation, we end up with minus the gradient of the pressure is the sum of 2 terms, one of which is $1+i \times U_z \frac{2\eta}{R^2 \rho_0 \omega}$, and that term is small compared to the first term and, typically, this says that the inertial term, $i \omega \rho_0 Z$, is large compared to the viscous drag term.

In this regime inertia is dominating and viscosity is not and it is because the shear-wave number is so small. The same approximations can be done for the effective density.

If we go to the other limit, take the case where S is less than 1, in this case $\eta/\omega \rho_0$ is large compared to R^2 (this is the low-frequency approximation), and now the viscous penetration depth is large, or occupies a significant fraction of the tube wall.

When we rewrite Newton's second law now, we end up with the other term being the larger term; that is, the viscous drag term is very large compared to the inertial term, so in this case we say that the flow is viscous-friction- dominated. You can make the same approximation for the effective modulus -- this result occurs -- and you can continue to do this for slits and I will not.

What have we done so far? Only for a single tube we write the shear-wave number in terms of the tube radius. Remember, we are trying to get to the ground and we are not going to measure the radius of tubes in the ground. We have looked at high- and low-frequency results and we can measure properties of bulk porous materials.

We can measure the flow resistivity and the tortuosity and the porosity of the bulk material. The goal now is to incorporate these numbers into bulk expressions for this effective density and

this effective modulus. Once we have that, we can then write down what the wave number is. As you recall from an earlier slide, I wrote that the wave number was equal to or proportional to the effective density divided by the effective modulus to the $\frac{1}{2}$, the plane-wave number of the plane wave propagating down the tube.

DR. MATULA: What is the physical significance of the difference between the 2 models?

DR. SABATIER: I will show that in just a minute. They are 2 different force shifts, that is the difference, so now we are going to say that everything can be scaled between those 2 is the idea.

DR. HOFER: Porous meaning being a bunch of pebbles tied together or spheres, then the curvature of the surfaces can be the opposite of what a tube would have, so it seems to me that that falls outside of the range between a tube and a parallel plane.

DR. SABATIER: It may. It turns out that packs of spheres behave quite nicely like tubes experimentally, at least when they are air-filled, when it is filled with air.

DR. ATCHLEY: In what frequency range?

DR. SABATIER: Up to the audio range.

DR. ATCHLEY: Penetration depth in terms of coordinate radius?

DR. SABATIER: What we will see is that to acoustically determine the flow resistivity or to acoustically determine the tortuosity we will have to take advantage of whether it is viscous-dominated or inertial dominated, so we will have to incorporate those effects in interpreting the data, and we will look at that in just a minute.

Right now we have demonstrated, for large S , it is inertial dominated, but we do not know what the physical parameter in the bulk material has to do with inertia. When it is viscous-dominated, it is the viscosity that is important but we do not know what bulk property -- I mean, it is going to be either tortuosity or flow resistivity, so you had better figure it out for yourself, and you probably already know it.

That is what I am going to do in the next part, but what I would like to do now is get a drink of water and take a 10-minute break.

DR. SABATIER: I did not set out to confuse you. I really set out to try to make this clear and go through in what I thought was a nice logical way to build up to flow resistivity and tortuosity in bulk materials.

There is a plane wave going down this tube, just a round tube like this, we are sending a plane wave. We had this boundary condition at the wall and the boundary condition at the wall forces the particle velocity to go to 0 at the wall.

The viscous penetration depth, which depends upon frequency, determines what the velocity profile across the tube looks like. It determines whether, if you wish, you get a plug (as I think Pat described it as) going down the tube or whether you have this big velocity parabolic profile going down the tube. The Bessel functions will describe what that profile looks like, but we are not particularly interested in that. I looked at 2 limits in Newton's second law.

[Transparency 27]

Now what I want to do is to go from a single tube to a whole bunch of tubes. Here is a stack of a porous material and we have n pores of radius R per unit cross-sectional area of tube. We have some average flow in the pore. We are going to assume all the pores are identical, to start with, and then we will figure out how we scale to different kinds of pores.

This looks like the way we started in the beginning, except we are now talking about what the pores might look like. We write the definition of the flow resistivity as $\Delta P, P_2 - P_1$, divided by $U \times d$, and you can replace $U \times U$ average, the velocity in the pore $\times N \pi R^2 \times d$.

For the D.C. flow, to go back now to Newton's second law and I take the case where ω was equal to 0, that is, the low-frequency approximation for the Newton's second law term, we realize from that case that the particle velocity down the tube is simply $R^2/8 \eta$ minus partial P with respect to Z -- I will go dig it up. This is an important step right here.

We neglect the first term on the right-hand side of the Newton's second law equation, so we end up with negative gradient of the pressure is equal to $8 \eta/R^2$ times the average velocity down the tube. We use that at $\omega=0$ to incorporate flow resistivity in the tube, so we substitute that into the expression for σ and you just do a little bit of algebra and you end up with S at the bottom of the page being equal to $8 \omega \rho_0/\sigma \Omega$.

[Transparency 28]

Now we have expressed S in terms of 2 measurable properties: the flow resistivity and the porosity, and we have gotten rid of R^2 and η , which is what I wanted to do. Now I can use Rudnick's apparatus in terms of trying to understand what is going on, or I can use that to get some numbers.

We do the same thing for a slit, starting right here. For a slit you end up with σ being $3 \eta/\omega A^2$, and do the same thing as I did before, and you can now write S for slits (S' , that is) as $3 \omega \rho_0/\sigma \Omega$.

[Transparencies 29-30]

Now I want to get tortuosity into that expression, so now, instead of having the pores running up and down, we let them be inclined at an angle θ . Now the porosity becomes $N \pi R^2$ divided by cosine θ for cylinders (you can work out the same thing for slits).

We plug in, then, the pressure gradient of the pore space. The sample length is D , but the pore length is L , so we end up with the pressure gradient in the pores being $P_2 - P_1/\cos \theta$, plug all that in, use $\omega=0$, again (the low-frequency approximation), and end up with the bottom expression there that σ is equal to $8 \eta/N \pi R^2$ to the fourth cosine θ and we substitute for ω being $N \pi R^2/\cos \theta$ -- and you get that expression for σ , $8 \eta/\Omega R^2 \cos^2 \theta$. You substitute again for η/R^2 and you get the expression I have in brackets now (in the box), which now gives us S in terms of all the measurable quantities. We can measure Q^2 , σ , and ω , and Q^2 is defined to be $1/\cos^2 \theta$, and the Q^2 is defined to be the tortuosity.

MR. SMITH: Is this random organization or --

DR. SABATIER: Just all the tubes going either this way or that way; they are all inclined to θ .

There are other notations for tortuosity and some of them are quite disheartening. The worst one is to go to Karman, which is flow of gases through porous materials (out of print), but he writes down about 5 expressions for tortuosity and he writes them down in terms of grain shapes -- all empirical. He writes one down in terms of porosity, which does not make any sense to me, so I avoid all of those, so get rid of that notation.

This notation is okay, Attenborough, I think, put this one forth in one of his papers and we will see, at least in looking at the acoustics in tortuosity, this one come back, again.

[Transparency 31]

The scaling between cylinders and slits was accomplished by two people. First it was done by Biot. Biot, in his original paper in 1950 or so, did this scaling. I showed what the complex density looked like for cylinders and slits and I showed what the effective bulk modulus looked like for cylinders and slits.

Then Attenborough did not quite do it right, but Attenborough made a scaling and it was scaled a little bit differently and a lot of people have complained about it. Attenborough showed that the complex density ρ is approximately equal to ρ_0 for slits with $S=5/3$.

Biot did it a little bit differently. Biot actually wrote something that he called the viscous correction factor, G_s , and it looks very similar (in fact, it can be shown to be the same in some sense, I think). There is a problem and I have never been able to figure it out.

Biot showed that he could scale between cylinders and slits if he let $S=4/3 S'$. There is a little difference in that one of them is scaling -- they are scaling 2 different equations, they are actually scaling 2 different sets of equations when they do it, and that is the difference.

[Transparency 32]

Here are Attenborough's results. Attenborough shows them both on the same page now. This is for frequencies between 100 Hz and 100 kHz. S (he actually calls it λ_d) is going between 0.1 and 3.5, so he simply empirically determines the scaling result.

[Transparency 33]

There is his result with $5/3$. He comments that, well, I cannot scale both the real and the imaginary parts, I can scale only the imaginary parts -- the real part does not quite scale, they are not the same.

He argues but, well, look at the size of the 2 numbers. One of them is all the way up to 1600 and the other one goes to only 15, so I do not have to worry about that detail. That is what Attenborough says and he has taken a lot of heat over this. It turns out that when you actually make the comparison to data it is hard to see this effect, this difference.

[Transparency 34]

Biot did it right, I think. He scales the viscous correction factor and his are close, but they are still not quite the same, so I think this discrepancy is probably what accounts for the argument that seems to have gone on between these two guys. Actually, Biot and Attenborough never argued. What happened was Arnott, Stenson, and some others suggested how it ought to be done.

When you look at the results and the wave numbers, the acoustic wave number, the plane-wave propagation number that you are going to try to measure, the phase speed and attenuation, you cannot see this effect.

[Transparency 35]

Keith Attenborough writes, then, that we will define a new shear wave number, S , to be $1/SP$ x all the stuff in the past, where for cylinders $SP=1$, and everything else is somewhere in-between. In his case, $SP=5/3$ for slits. It is my experience, then, that you do not need to develop pore-shape scaling for all other kinds of other kinds of shapes, at least not for the level of accuracy of measurements that I can make. It may be that if you get away from soils and start looking at these kinds of tubes, that may be what you can do, but I have not been able to tell the effect.

[Transparency 36]

I summarize, then, in this low-frequency approximation that we have 4 properties that we have to worry about, ω , Q^2 sigma, pore shape. The bulk wave number that is now ω/C , which is frequency over the wave speed that we are going to measure, is ω times this effective density divided by the effective compressibility. If we plug in the numbers for ρ and for the effective density, we can show that then K_b is equal to this expression, where K_b is the bulk wave number in the material, K_0 is the wave number in air, ω/C for air, γ is the ratio-specific heat, A is the constant -- it is about equal to 1 -- Q^2 is the tortuosity, and then we have this imaginary term $L(1/2)$.

If you look at the high-frequency low resistivity, K_b is equal to $K_0 Q A^2 (1/2)$. What I say there is that γA is almost 1. If γA is approximately equal to 1, C_0/C_B is approximately equal to the square root of the tortuosity. It is not quite the acoustic index of refraction; it is only approximately equal to the acoustic index of refraction. It depends on the product of the $A \gamma$. Prandtl number is $3/4$ and γ is 1.4 and you do that calculation.

[Transparency 37]

Now we consider further low-frequency approximation of this low-frequency approximation for the bulk wave number. What I wanted to consider was to just look at the wave number K_b . If I look at it, remember that we can always write a wave number -- and this is the wave number I measured when I measured the attenuation and the phase velocity with the probe microphone buried in sand -- we can always write it as a real part plus i times some imaginary part.

What we would like to do, then, is to square this thing, or square that thing on this side over here, and then consider the real part of this term squared. The real part of this squared is KR^2 minus KI^2 , the real part of that squared, and that is approximately equal to something like $AQ^2 \omega^2$, where ω now is frequency.

The imaginary part of K_b^2 is $2 KR KI$, so that is now equal to γ frequency over sound speed squared, flow resistivity, porosity divided by $\rho_0 \omega$, which is proportional to ω to the first power. So we now have 2 frequency dependencies that we can now verify in our measurements.

We would like to experimentally consider this frequency dependence or we might like to measure KR and KI for a porous material and then invert the data to yield values of these 4 parameters.

[Transparency 38]

Here is what I did when I pushed the probe in the ground. This is the loudspeaker (and let's assume it is at normal incidence) and we measure, then, the probe response to the loudspeaker at 2 different positions, D_1 and D_2 . We write down the transfer function between those 2; that is, we divide the signal between those 2 points.

The phase between these 2 points is the arc tangent of the imaginary part over the real part, which is $KR \times \Delta D$, and we have to worry about 2π phase flips, so you have to be careful when you first start doing this. You have to make sure that this distance does not correspond to more than 2π or you will lose the 2π nature of these terms right here.

The magnitude of this transfer function in decibels is just 10 times the log base 10 of the magnitude of T^2 or $20 \log_{10} E$, $E(-KI \Delta D)$, or this expression right here, so now we can

write down KR is this phase that we measured here divided by ΔD or KI , the imaginary part of the number, is just this expression right here. Now we can write down, if you wish, these 2 numbers, which we can put back together to construct the bulk wave number in terms of the model.

[Transparency 39]

I will show you what the microphone looks like just for a minute, because you have to be careful when you design this microphone. You have to protect the end of the microphone, so what you usually do is you put it in a nose cone, and you have to drill some holes in the nose cone.

As soon as you put holes in this nose cone, you make a Helmholtz resonator, so you have to worry about the frequency response to this Helmholtz resonator. You can design these. In the next slide I show the frequency response of this probe microphone.

The way you operate this thing is you put the sensing end in the ground and the non-sensing end is sticking up in the air where the loudspeaker is. There is a lot more sound in the air than there is in the ground and you want to be very careful to make sure that you do not get sound coming down through the tube. I have never had a problem with it, but since the microphone is bigger than the cable coming out, I just filled it with glass beads.

[Transparency 40]

This is the response, then, of this probe microphone. We actually designed this probe microphone to have a resonance at about 5500 Hz and I did not bring those data, but here is the frequency response of the probe microphone measured in air and this is the resonance.

When you stick this in soil, those little holes, those 24 holes we had drilled in there, can plug up with dirt, so that will then change the frequency response of this probe microphone, so you have to worry about that as well.

Here it is, then, with 0 holes filled -- that is these data out here. Then we fill 19 of the holes and it is this response. Twenty-three holes and you get that response. You fill all the holes and you get that response. What do you fill them with? We filled them with Play-Do, clay from Walmart. Everything is porous, that is one thing I learned a long time ago, concrete is porous, wood is porous, asphalt is porous, nothing has absolute impedance, as you might think.

Whenever you take the tube out of the ground and you look at this probe microphone, it is important to make sure that you have the holes open for the frequency range you are trying to work in, but it is pretty forgiving, and we had it designed that way. Initially, when we did not design it that way, we found out that the resonance was moving way down all the time and giving us lots of problems.

[Transparency 41]

This is about as careful data as we have been able to get in the laboratory. It is 500 μm spherical glass beads, the same kind of experimental configuration as we have seen before, and we are displaying the magnitude and the phase difference as a function of depth (this is just transfer function) from 1 cm to 8 cm.

This obviously is a log scale and we clearly see that there are some places where the difference in the magnitude at a particular frequency as a function of depth is not very uniform. It has more to do with our experimental apparatus than it does anything else, and how we measure the depth, and this is invasive. When you push this thing in, the glass beads move

around. They have to go somewhere and they bulge up on the surface and they change the surface a little bit.

[Transparency 42]

This is the imaginary part of K_b^2 and we remember it went as ω to the first power for spherical glass beads of a diameter that I do not seem to know. It is linear. It wiggles, but it is generally linear, and the line just happens to be a linear fit to the first power of ω .

[Transparency 43]

This is the real part of K_b^2 and we remember that that went as ω^2 in our simple model, and there is what the data look like for that case.

DR. MOLDOVER: This is a moderate porous material?

DR. SABATIER: Yes, spherical glass beads, just a big vat of spherical glass beads, and the measurement is made outdoors. You could do this indoors, but when you start going to low frequencies, you start having problems with lots of room modes. Remember, there is a microphone at the surface and we are trying to do a transfer function, so if you get a lot of room modes in this, you can contaminate the experiment very badly, so you dig a hole in the ground and you dump all these beads in the ground and you make measurements.

We have not designed an experiment that is supposed to be -- the depth of the probe is probably accurate to only a millimeter and we are moving it down centimeters at a time. We are trying to understand what we are going to do outdoors, or what we have done outdoors, with the probe microphone.

[Transparency 44]

This shows, inverting those data for flow resistivity versus flow resistivity measured with Leonard's apparatus, for five spherical glass bead sizes that range from 60 μ to 275 μ , the lower axis is the Leonard's apparatus flow resistivity measurement and, on the vertical axis, is the acoustic flow resistivity measurement.

If we go back and look at the expression for K_b^2 , there was $SP \sigma \omega$, so we have set SP equal to 1 (assume it is cylinders). We can measure that a random packing of spherical glass beads has a porosity of 0.38, and you can either do that yourself or find it in the literature.

The vertical axis, then, is $\sigma \omega$ and this is σ and the slope of the line is 0.38, the porosity, which is what we expect. This is the best that I have been able to do to show that I can measure the bulk flow resistivity of a porous material with a microphone.

PARTICIPANT: Does that depend on the size of the glass beads?

DR. SABATIER: Yes, so you see that the flow resistivity changes for sizes of glass beads.

DR. GARRETT: I would like to point out that the triangle up there is a bad choice of symbol. I did not realize that was the last data point, that the data actually run all the way to --

DR. SABATIER: I apologize. That is a very important data point and we did not do that data point initially. We are going over maybe 2 orders of magnitude in flow resistivity, so over 2 orders of magnitude we have this linear fit.

MR. POULSEN: Those points, 275, and the next one is 125, and then you have 190.

DR. SABATIER: There may be a mistake in the numbers, there may be an error here. We go from 60 μ all the way up to 275 μ and everything else is in-between, and it goes sequentially, they are not out of order. Forget that, I see what you are telling me. Yes, there is a mistake. I apologize for that.

This was a slide that was done by an undergraduate student who came in and wanted something to do, a civil engineering major, and you get a lot of civil engineering majors who are interested in this kind of stuff, so I said, here, go take 5 glasses of sized beads and do this, make these measurements, so this was done by a student, an undergraduate student and presented at one of our regional state meetings.

People have looked at fractal distributions of pore sizes and have developed the wave numbers for bimodal distributions of glass beads. There is all kinds of that stuff in the literature.

DR. KEOLIAN: But always homogeneous?

DR. SABATIER: But always homogeneous, that is exactly right.

The next thing is to do something like Hank Bass' measurement, the initial one I talked about. Yes?

DR. KEOLIAN: I have trouble understanding how you can get away with pushing this microphone into the dirt, like, for instance, leaving a little air hole around the probe that will --

DR. SABATIER: What you do is, and I did not point this out, you design the nose cone of the microphone to have a smaller diameter than the body. Then when you go to soils, you have to construct holes that are smaller than the size of the tube. You have a jig, push it in, and then just before the bottom, when you get to the depth where you stop pushing it in, you squirt a little bead of STP oil treatment around and that is the seal, and it works nicely.

DR. KEOLIAN: So where you shove the probe --

DR. SABATIER: You make a hole first.

DR. WAKELAND: Is it drilling or are you --

DR. SABATIER: It depends on the soil. Civil engineers do this all the time. There are some soils where you can take a cylindrical tube with a knife edge on it and you put in the jig, just tap it in, and it comes out and makes a perfectly clean plug, makes a beautiful hole.

There are some when you pull it out it excavates the soil, so for that one you have to have a drill. All these kinds of things are available within the soil science community. Do not try to do it the way Hank Bass says you do. Hank Bass says you pound it in with a hammer. That is not what you do, that does not work. If you do do that, you lose the seal and the seal is very important.

DR. HOFER: Does it matter if there is a little air gap around the nose cone?

DR. SABATIER: No, we cannot see that effect. In fact, what we have done is we put together kinds of engineering devices for people -- the agricultural community has a lot of interest in measuring permeability or flow resistivity of soils -- so all they want is a device that gives them a number, so we put a computer in, a bunch of buttons, that tells you what to do.

We used to drill a 1-cm hole and then a 2-cm hole, and pretty soon you cannot use that same hole, so they came up with the idea that we want to make just one hole, a 10-cm hole, and make measurements every centimeter and push it down, so we could not see the effect of that if you had that big hole below the mike.

[Transparency 45]

Back to Henry's measurement. In the literature it is typically called Lloyd's mirror experiment or the image problem in electromagnetic theory and we have a source above the ground, a receiver of some height or one at a distance, or 2, an angle between the incident ray in the ground and we can write the total pressure at R -- over here. Then you have to go to either

Sabi Hayat at Penn State and Keith Attenborough and several other people before them who worked out -- this is the point source, so this has to be the spherical wave reflection coefficient -- the approximation for this.

I do not want to worry about it. What I want to tell you is just that Q is a function of Z , the impedance, and we can write Z in terms of this wave number, and the density, so we can get Z . This impedance now depends upon tortuosity, σ , and flow resistivity.

When you do this experiment, the pressure here is the sum of a direct wave and a reflected wave, so you get interference every time $R_1 - R_2$ corresponds to a half-wavelength and you get this series of minima in this expression.

This is strongly dependent on geometry and subtly dependent on Z on the surface impedance. In fact, it is also dependent upon the roughness of the ground. If the ground has some roughness, then you begin to not get -- you get reflection in other directions, so now you have to start worrying about the effects of roughness, so they end up putting another parameter in here, in this height distribution, the number of roughness elements per unit area.

We try to make measurements on flat porous surfaces of spherical beads to verify this.

MR. WAKELAND: Can you really see this at smaller than a wavelength?

DR. SABATIER: Yes, of course you can see the roughness. In fact, you have to worry about whether or not -- depending on the theory that you use to understand the roughness, model the roughness, you can do Rayleigh scattering.

MR. WAKELAND: I mean the surface of the glass beads is rough.

DR. SABATIER: No, we do not see the glass bead roughness. Michael Oelze, in the back of the room, is developing a prospectus to take these notions and apply this to a back scatter experiment looking at roughness to try to back out what roughness is.

We can write, then, using our low-frequency approximation Z as the normalized impedance, $1/\rho_0 C_0 \times \omega$ times the root of the effective density over the bulk wave number K_b , and the low-frequency approximation we end up with $C = 1 + i \sigma / \omega$.

Two points to make: One is that the real and imaginary parts are equal in this model, and when I initially showed Hank's impedance data, I commented that the real and imaginary parts were about the same, so that comes out in this approximation. The other one is that the high- and low-frequency approximations for Z depend upon σ / ω or they depend upon Q^2 / ω^2 .

A whole lot of effort was spent by several people using this model, this Lloyd's mirror model, to back out all the parameters, the flow resistivity, the porosity, the tortuosity, and the pore shape. They were trying to invert the data and they used all kinds of geometries.

What you do is you measure this frequency response for a lot of geometries, for a lot of ranges, collect a whole lot of data over the same spot, same area of the ground, and then do a least squares fit to those data of all the data, do a simulated annealing, to come up with what the parameters are, and they would come out with all 4 parameters.

What we are showing right here is that you cannot do that. It does not work, but a lot of work was done and it has all been straightened out now. This notion is now being developed into a standard by the ASA for measuring the impedance of grounds. It is very important to measure the impedance of grounds.

There is a whole lot of concern about community noise, and if you are going to design terrains for community noise, you need to know the ground impedance. Dagel and others who

worry about outdoor sound propagation -- Richard Raspit and Mike White have done this for a couple of years, John Noble, who is one of our former graduate students, have developed a standard based on the Lloyd's mirror measurement and this model of the ground (they have a couple of models, actually), they developed a rigid-backed layer and the only difference we get when you do a rigid-backed layer is that Z is equal to this normalization times i times the cotangent of $K_b D$, where K_b is the wave number and D is the thickness of the layer.

DR. HARGROVE: How high is that number?

DR. SABATIER: Ten kHz. Ten kHz is a very high frequency for us. Five kHz is high enough for everything we are doing.

MR. WAKELAND: And at 10 kHz the underlying model you think is still good in terms of the rigid --

DR. SABATIER: Yes, this is all rigid.

MR. WAKELAND: So you say 10 kHz is high, is 5 kHz high? I mean, where is --

DR. SABATIER: We can get this high-frequency limit of the tortuosity out at 5 kHz. High means that the imaginary part of the wave number ends up being large compared to the real part -- maybe I have that backward, so you can neglect it. One percent.

[Transparency 46]

Carl Fredrickson now has calculated something called level difference, which I should explain. I showed the Lloyd's mirror experiment, where you had a source and a receiver. That is a hard experiment to do, because you do not know the reference sound pressure level, so you have to compute a transfer function.

You put another microphone right here, so you have 2 microphones now. You take the transfer function between these 2 microphones and the idea is to take out the source. It makes it a little more complicated but it is all done on a computer anyway, so it does not matter.

This is the ratio, the sound pressure level difference in dB for a particular geometry, which is probably spelled out down here. The source height is 40 cm, the range is 1.63 m, the top microphone is 40 cm, and the bottom microphone is 10 cm, and sets the bulk flow resistivity equal to 75 kilorayles, Q^2 at 1.9, and porosity equal to 0.37.

He then uses the expressions for K_b to develop the impedance for cylinders, slits, and triangles, so we conclude that we do not need to worry about pore shape or different pore shapes from this.

These are some data, and the calculation. Carl Fredrickson did this. Those are the kinds of effects we get. These deviations that you see are due -- probably they are due to the lack of a point source. We work very hard at trying to make a point source, but you have to mount -- also, they could be due to mounts that you support the microphones with. You work very hard at trying to make all these things go away and this is about as good as we can get in making things go away. Appreciate that the deviations there are only about a dB, so we are talking about only a 10% error.

This is a good time for a break. When we come back, what I would like to talk about is coupling into the sound in the ground when the ground is poro-elastic.

DR. SABATIER: We are going to now talk about coupling of sound into a poro-elastic material as opposed to a rigid porous material, so we are now going to allow the matrix to have an elasticity.

[Transparency 47]

We have two sensors, so we have this probe microphone sensor that we can use to measure pore pressure and now you can use a geophone to measure the matrix frame velocity. If you want to look at the elasticity, then the rigid frame model is not sufficient any more; you have to couple it with an elastic model.

This is what Maurice Biot did in 1956 while he was at Exxon (not the same Biot). It has been used extensively in the rock physics community, but I think they might be well beyond Biot's model now. It was used extensively by people in the ocean sediments community and there is probably nobody else who worries about the elasticity of the ground for air-coupled sounds than myself and Hank Bass. (I may be wrong, though.)

[Transparency 48]

I will speed up a bit, because I have gone through only half of the material I wanted to go through. Biot used the Lagrangian formalism to derive the equations. In poro-elastic material the grain sizes and the pore sizes are small compared to a wavelength. The measurement volume or the idealized volume that he used to develop the equations is large so that it represents many grains and many pores in the volume that we are talking about.

These are his equations. They are not quite Biot's equations, these are Biot's equations written in a form by Stole, who used these extensively in trying to talk about properties of ocean sediments. In these equations, H , C , and M are elastic constants, which I will come to in a minute, ρ and ρ fluid are as before, the frame and the fluid density. The term M that sits over here in front of MC times the second partial with respect to T^2 is $Q^2 \rho$ fluid/ ω , which is in the rigid frame theory, σ is the bulk flow resistivity.

Biot did not do flow resistivity, Biot did permeability, but I only know how to measure flow resistivity, so you just convert it. Biot defines this to be his dynamic permeability -- in my case, dynamic flow resistivity -- so this is the viscous correction factor, or this is the expression that we derive for as cylindrical tube or for a parallel slit.

S is as defined before, where it has the properties of interest, Q^2 , σ and ω . All of the scaling that has already been done Biot did when he developed this model, so he did all of that scaling that I went through for rigid frame tubes. In Biot theory the tubes are rigid and the elasticity comes in through the elasticity of the matrix.

[Transparency 49]

Those are complicated equations, but just for completeness the main thing I want to point out right here is what C is in these equations. Little u is defined to be the displacement of the matrix, U is the displacement of the fluid, so Stole wanted to talk about this relative displacement here, $u-U$, in his work.

The elastic constants H , C , and M are derived and end up looking like this, H , C , M , and D , I guess, and K_b is not the wave number. K_b is defined to be the frame modulus of the bulk material, K_f is the grain modulus, that is, the modulus of a single crystal of the composite mineral material that the grains are made of. K_f is the fluid bulk modulus, but Biot, since he was working with water-filled porous materials, did not need to worry about thermal conduction, so he did not include temperature effects.

He suggested in his original paper some Gedanken experiments for these and, in particular, he called them "jacketed" and "unjacketed" experiments. In the jacketed experiments you

conceptually put this porous material in a bag and the bag has an opening, so the pore pressure can be relieved as you try to measure the bulk modulus. In the other case you do the standard kind of compression modulus.

I have never measured these things. In fact, there seems to be a controversy in the literature in underwater sediment. Nobody has ever measured them, but stole everybody what they were. A guy at Texas, Nick Chitoris, claims that he has measured the slow wave in water-filled sediments -- actually, it is sand that they have in a big tank right there.

He had to end up with a grain modulus and, I guess, a bulk modulus that was much different from what has been used in the literature. He has actually measured it and people do not like his measurements, but no one else, apparently, has measured this thing. I will talk about other ways of getting these for air-filled soils.

[Transparency 50]

If you are going to look at the coupling of sound in the ground, then you have got to write down the stress-strain relationships and here they are in Biot notation, where you include both the fluid and the solid, and I just put them here for reference. I will look at some limits of these that are probably important.

Suppose that the medium is a fluid. The porosity goes to 1. Then this expression for the pore fluid pressure just reduces to the fluid bulk modulus times the fluid displacement. If it is an elastic solid, the porosity goes to 0 and this elastic constant H just reduces to what we get in the normal elastic material where we have no porosity, the bulk modulus minus $4/3$ of the shear modulus.

You can also take Biot's equations and reduce them to 2 limits. You can reduce them to the rigid frame limit and in the rigid frame limit of Biot's equations you get this equation right here, which would look just like the effective medium porous frame rigid wave equation if I had written it down, but I did not. This is, then, the wave equation in elastic solids. These are all for compressional waves. Biot wrote down equations for shear waves as well.

[Transparency 51]

The interesting thing that comes out of Biot's theory, of course, is that there are 2 wave types. If you take 2 solutions, 2 plane wave solutions, one for E and one for C , substitute those into the wave equations and then write down that the determinant of the coefficients has to be 0, and solve those, you find your solution is quadratic and L^2 and you end up with 2 wave types. Biot referred to these as type I and type II waves.

These 2 waves have dramatically different characteristics. The type I wave is typically nondispersive, has little attenuation. The type II wave is dispersive and has a lot of attenuation. It has a much slower wave speed. In the literature they have been called the fast and the slow waves.

In porous materials you can find that the 2 wave types switch speeds, the slow wave can become faster than the fast wave, so it seems inappropriate to use those words, but people do all the time.

You can also solve for the ratio of the amplitudes, that is, what is the ratio of the displacements for these 2 wave types. We have 2 waves moving at 2 different speeds, two different propagation wave numbers, and we can arrive at what their displacements are. Of course, you have to solve this expression for the various wave numbers that you have there.

Plona at Schlumberger-Dole was able to take a fused glass bead sample up around a megahertz to be able to measure the slow wave in this water-filled material. Even in his measurement it was necessary to look at the critical angle; that is, to turn the sample until we get over here, so that the fast wave, the type I wave, which is faster than the speed of sound in water, gets to a critical angle and then it is damped evanescently and then you can see the slow wave come out.

The same thing was done by Nick Chitoris in sand at the University of Texas. It is still controversial; no one else has ever seen it. There is another whole group that says it is something else.

[Transparency 52]

These are some results of looking at the phase velocity and the attenuation for the 2 wave types that come out of Biot's theory. This is for air-filled soils, so we are putting in a flow resistivity, a porosity, and a tortuosity that are reasonable numbers for soils.

I have not said it yet, but we can measure the shear wave speed in soils and we can measure the compression wave speed in soils. If we do those 2, we can then go in and back out what the shear modulus is, what the frame modulus is, and to get the grain modulus we do a calculation of what the mineral composition is of the soil to get the grain modulus, so we do not measure these parameters as maybe we ought to be doing.

This is showing just the wave speed and the attenuation. What you see in these slides is that the frame velocity is 270 m/sec with very little dispersion between 0 and 2.5 kHz. This is the frequency range we operate in for our measurements, and you see what the attenuation coefficient is. It is quite small. That is the type I wave at the top.

For the type II wave we get something that is essentially the same results we do if we use the rigid frame theory. We do not have to use Biot's theory to get this result. We get a large dispersion and a lot of attenuation. Notice the attenuation is orders of magnitude higher for this wave compared to that wave, and that the wave speed is quite dispersive; that is, you go through this rather dramatic change in wave speed.

[Transparency 53]

Soils are a special case of Biot's theory. This shows the ratio of the amplitudes of the fluid and solid displacement for the 2 types of waves. For the type I wave at the top we are writing the magnitude of the displacement between the fluid and the solid for the type I wave and we see that it goes from 1 to about 0.7 of the frequency range of interest and that the phase goes through whatever it goes through, 0.25 to 0.05.

If we look at the same results for the slow wave, though, we see that the relative amplitude is about 1×10^{-3} , so the relative amplitude, then, is 1000 times smaller and the phase is 1 as opposed to something that is a lot smaller.

Normally what people say is that the solid and the fluid are moving in phase for the type I wave and for the type II wave the solid and fluid are moving out of phase.

MR. HODGSON: Are these experimental results or just calculations?

DR. SABATIER: Just calculations.

[Transparency 54]

If we are going to look at coupling of sound into the ground, then we have to have a set of boundary conditions and we have to have different kinds of boundary conditions, at least for

different kinds of boundaries. The ground first starts off with an air-porous material interface and at some point in our models we have allowed the -- the ground is highly porous in the first foot but porosity gets small right away. The flow resistivity goes up very high as you go down.

We have modeled the coupling of sound to the ground as an air-porous boundary, porous-solid boundary, over a solid-nonporous boundary, so an elastic material with a poro-elastic substrate topped by air, so a sandwich.

There are 6 boundary conditions, in general: continuity of normal stress; normal velocity; tangential stress; tangential velocity; porosity, average pore pressure; and normal component of the fluid velocity. They are all continuous.

For a porous-solid over a solid surface, one of the main things to note here is that this one goes to 0, the difference in the velocities between the 2 at this nonporous interface.

A similar set of boundary conditions has to be written down at the fluid or the atmosphere porous soil surface and these are the boundary conditions that you write down.

For this simple model, then, we end up with 9 boundary conditions that we have to solve with Biot's equations.

[Transparency 55]

Before I go any further, though, I want to show some results of some measurements taken by Craig Hickey in the laboratory that experimentally highlights some of the properties of these type II waves, we think.

This is a vat of sand that is a meter in diameter by about a meter deep. It is air-filled. In this vat of sand we have installed microphones and geophones, both horizontal- and vertical-component geophones, at a host of depths. The depths are indicated by the density of the measurements on those lines right there.

Zero corresponds to the top of the tank and 50 cm corresponds to the depth of the deepest sensor. The color code tells you the kinds of sensors that we have. Down at 45 or 50 cm there are some vertical-component geophones in this graph. We have also buried some microphones down in the porous material and that goes all the way up to near the surface, to 5 cm.

The squares are the positions where we made probe measurements as a function of depth. The experimental configuration is that the sound source -- this is the tank and we have a loudspeaker up here and the sensors are down in this tank.

PARTICIPANT: What is the difference between the microphones and the geophones?

DR. SABATIER: The displacements of the solid are very small, so we believe that the microphone measures the response of the fluid pressure. The geophone is surrounded by the matrix of the soil and is more efficiently impedance-matched to the matrix, so we believe that it measures the solid motion. That is our assumption.

These are time-of-flight measurements versus depth for these sensors. In the shallow depths all the sensors give us the same numbers. Then we go through a transition region that we do not believe is a layer depth, we do not believe there is a discontinuity in the layering of the soil -- we have very carefully tried to construct this soil layer so it does not have layers -- and then we go through this transition region, at which time we end up with another slope where the microphones and the geophones give us the same velocities.

That is for a frequency probably of 1 kHz.

[Transparency 56]

We repeat the experiment but instead of using a loudspeaker as the source we use the mechanical shaker, a small mechanical shaker with a disk sitting on the surface, and we drive that with a mechanical shaker and repeat the experiment.

For 4 different frequencies we show the velocity that we get with geophones or microphones from 5 cm to 50 cm in increments of every 5 cm. The results are that we see no dispersion in these data. I do not have to convince you or show you that at the shallow depth with the microphone we are going to get a lot of dispersion. We already know that, we have done that, we have demonstrated that.

If there is a mechanical shaker on the surface, we get a wave that has no dispersion. If there is a loudspeaker in the air, we get 2 different wave speeds with depth. The shallow wave speed is very dispersive and the other one is not. I think between what I have showed you in this slide and the previous slide that you can come to that conclusion.

[Transparency 57 – unavailable at time of printing]

One other measurement is to look at the magnitude of the microphone voltages as a function of depth and this is the probe microphone magnitude as a function of the depth. We have a loudspeaker in the air and we are looking at the magnitude of the voltage out of the probe microphone from 0 out to about 20 cm, down to 20 cm in depth, and we get this, we get 2 slopes. This was at 2 kHz. There is a transition region for this slope that I do not show right here, just as there was a transitional region for the velocity. This is the understanding of what is going on.

When you have the loudspeaker in the air and you drive the ground, you impedance match very efficiently to the pore fluid and you excite the rigid frame wave (that is the Biot flow wave). It is very dispersive. It is highly attenuated. You do not put a lot of energy into the matrix wave, to the solid-borne wave, because there is a big impedance mismatch between the air and the solid.

As you go down in depth, the type II wave is damped away exponentially below the amplitude -- eventually, in the transition region, with an amplitude of approximately equal to the type I wave but, below that, the type I wave has a larger amplitude, and then you see the velocity of the type I wave and then you see its attenuation.

MR. WAKELAND: In this case the velocity of the slow wave is higher than the velocity of the type II?

DR. SABATIER: No, it is not. It is 140 m/sec compared to 207 m/sec. Let me put that back up.

This graph might have confused you. The velocity is not the slope. It is 1/slope.

MR. WAKELAND: Oh, it is the time. Thank you.

MR. HARDIMAN: But your geophone is measuring the same thing down there, right? Doesn't that contradict --

DR. SABATIER: We do not think so. We think what the geophone and the microphone are measuring is one is measuring the pressure of the type I wave and the other one is measuring the displacement of the type I wave. At the surface the geophone is measuring the displacement of the solid.

I think there is one thing that has to be done before we can prove this conclusively and that is to predict the pressure as a function of depth.

Yes, we had shear waves in here. We have shear transducers and there are some results on shear waves, and we do not see this effect that you see right here.

MR. HICKEY: Before you showed the second slide I would have said we were looking at the effects of compaction here.

DR. SABATIER: That is right and I want to say that is not what we are looking at, the effects of compaction.

There is a problem. There ought to be a glitch right here, there should be a glitch in the arrival times and we do not see that glitch. We are still trying to find that. We need to find the glitch, so I think that is one problem if we do not see that effect. Explain it, Craig.

MR. HICKEY: In the transition region, this is a pulse, so a pulse-type measurement, so in the transition region, where the amplitude of these 2 waves should be about the same, you should see 2 pulses there. You do see some distortion but we cannot separate these things in time.

Explanations of these measurements are kind of argued from the end, from the extremes. At very deep the pulse we see is associated with the fast wave and at very shallow the pulse is associated with the slow wave and in-between, where the energy of these 2 waves comes to be about the same, we do see a distortion but we cannot see 2 separate pulses.

MR. WAKELAND: So you are putting a pulse into the speaker, do the shapes look a lot different in the geophone?

MR. HICKEY: No, they do look similar. I mean, they are reversed in polarity but, other than that, there is no difference.

DR. SABATIER: We have a very good model for the pore-fluid pressure due to the slow wave, we think we have a very good model for that effect in a rigid frame porous material.

[Transparency 58]

What we have to do is calculate the pore pressure and the velocity associated with the slow wave. If we can really measure, if we really are measuring the velocity and the pressure of the slow wave and the fast wave, then we should be able to predict those 2.

Our problem is we do not know how to calibrate the probe microphone in the porous material. We can take only transfer function measurements, and Tom has to tell us how to answer that. I talked to Steve about reciprocity (he probably does not remember this) and he spent a long time trying to get me to understand how I was going to do reciprocity, so I came back and did what you told me, but I cannot get any signal out of my microphone to go any distance in the porous material.

I just took the end caps off of a microphone and used them as transmitters and receivers. A microphone is not a very good loudspeaker, and I need a microphone that is -- I need to use a microphone to do this in the porous material because I want to excite the slow wave.

When we put that mechanical shaker on the surface, we do not see slow waves. To see the slow wave we need to have the relative motion between the fluid and the solid. With a mechanical shaker we do not see these 2 wave types. We can see it only if we use a loudspeaker to excite the 2 wave types. That is what we need help with.

MR. HARRIS: You would excite a fast wave in the solid, so you would expect to see that wave with the loudspeaker.

DR. SABATIER: Yes, if we make the porosity go away.

We do excite both wave types with the loudspeaker, but we put more energy into one than the other.

MR. HARRIS: So maybe that is an observational difficulty. I mean, that wave may exist there.

DR. SABATIER: You are commenting about my shaker does not make the slow wave. I did not know what you were commenting about; I was already looking at my next slide.

Recall earlier that I had a geophone in the ground back with Hank Bass' experiment and we were measuring the frequency response of this vertical-component geophone. If we are now going to make a prediction of that, this is the problem, the boundary conditions with all the wave type problems that we have to solve.

We have 2 boundaries with 9 boundary conditions, an incident airborne wave, a reflected airborne wave, transmitted shear waves that have fluid and solid displacements, a type II wave, and a type I wave. Everything has fluid and solid displacements. This is nonporous, so there is no slow wave in the lower material, but there are both shear and compressional waves in the lower material.

We can measure some things.

DR. HAMILTON: Is there a slow and fast wave for the shear mode?

DR. SABATIER: No, it turns out not to be the case.

We can measure the acoustic surface impedance just in this kind of notation, which is only for numerical work. We can say that the surface pressure is minus i times the fluid modulus wave number times the sum of the incident surface waves, where i is 1 and 2 for fast and slow. The same thing for the surface velocity. The impedance is the ratio of those 2 numbers, so we can do that. That is something we can calculate from this.

We can calculate the pressure below the surface in the porous material and we have to sum up the pressures for all of the upgoing waves and for all the downgoing waves for the 2 types of waves, and the same thing for the normal component of velocity (or we can do the horizontal component of velocity), and we have to sum up all the components for the 2 types of waves that are going up and down. It is only numerical and it is all you can do, but some interesting results happen.

[Transparency 59]

That was a calculation of the surface impedance. What we have done is to make the lower material, the substrate material, have a very high reflection coefficient; that is, it has very high compressional wave speeds and very high shear wave speeds (10 times what we measured them to be).

If we do that, we get a lot of energy reflected back there and we get strong standing waves in this layer, and it shows up in the surface impedance. When we first tried to publish this the reviewers did not like it and they told us the numerical artifacts. The twin structures associated with shear and compression, compression and shear wave -- why you have twin structures. I very well understood it, but I never could convince the reviewers that I understood it, so I published it in a little obscure journal that no one would care about. I will come back to that.

If you get off axis, then these shear and compression modes can interfere, so you get weird things happening and the impedance can be everywhere. The reason I did this is because if you

go back to Hank Bass' data or you go look at anybody else's impedance data of the ground it is everywhere. I used to draw lines through their data -- I will go back and show you that.

[return to Transparency 2]

I mean, look at that. There is structure all over that thing, you do not have to stare at it very long. Then you can find some that have real structure. You can actually go look at impedance data where they have structures like I saw but they did not exist at the time we did these calculations.

The calculated line is an attempt to calculate the impedance using a semi-infinite rigid frame model, so only in the rigid frame theory what is effectively the slow wave, a semi-infinite Biot model, and you get the same results for the 2, which are indicating that you do not need the Biot model to predict the impedance if there is no structure.

If there is some structure in your impedance data, then you have to have this elastic model.

Snow shows resonances like this. The impedance of snow show these kinds of effects, but it is not due to the elastic properties. It is due to the wave number or the attenuation is low enough that the pore fluid wave gets all the way down and back and you end up with resonances. You can see this in pea gravel, you can see this in some sands, you can see this in some highly permeable soils.

[Transparency 60]

This was an attempt to make a calculation of the vertical geophone measures in this soil layer, so it is just looking at interferences of upgoing and downgoing waves in a porous material.

One of the problems that has not been solved is what is the attenuation of the soil material for the type II wave. All we have done here is to play the game that has been done in ocean sediments and that is to allow for a small imaginary part of the compression wave speed or the shear wave speed to keep these peaks from going too high. They go up about a factor of 2. At least we could make attempts to predict what the attenuation coefficient is, but we have not measured it, nor do we have a model.

In the rocks physics community there is not much progress as well in doing the same thing.

[Transparency 61]

Now I would like to talk about something that we are currently working on that relates to at least having a good understanding of porous materials. On the top of that forklift is a laser Doppler vibrometer, which is just an interferometer. You send out a laser beam that is modulated through a Bragg cell and you collect the reflected light from the surface of the ground and you demodulate that light and you can measure, then, the surface velocity of whatever you think the laser is bouncing off of.

My suggestion to you is that you are measuring the matrix velocity of the slow wave if you use a loudspeaker to excite the ground. What we do is take a loudspeaker and then insonify [phonetic] the ground and then read out the surface velocity of the slow wave using the laser.

[Transparency 62]

Here you can see the loudspeaker -- it is in very close right here. That is the horn right there. The laser is sitting off 3 m away. This is a 16-inch-by-16-inch spot on the ground -- there is some distortion, obviously. There is an aluminum disk 4 inches in diameter buried 2 cm below the surface.

We have made these measurements for a long period of time looking for things buried in the sand using geophones. That is, you buy from Mark Products the guts of the geophone, get rid of the case, so you end up with something that has a cylindrical base of about a centimeter in diameter, so it is going to average the velocity over a centimeter.

We measure exactly what I showed you on this previous curve. We measure the velocity versus frequency of the coupling of sound into the ground and it is the matrix displacement of the slow wave that we are measuring and we integrate it between some band -- you do not have to integrate, you look. You look at the signature where there is something buried and where there is not something buried, and there is from a factor of 2 to a factor of 10 difference, depending on the properties of the ground, and we can see that.

Now, with this laser, we have a spot size that is a fraction of a millimeter, let's say it is on the order of a millimeter, as opposed to something that averages over the size of the geophone base. We are insonifying the ground with an airborne source and we are reading out the velocity. This is a 16-inch grid and there are 64 by 64 laser measurements made, that is, surface particle velocity measurements made with the laser. The target is where those 4 red spots are there right there, so the resolution is on the order of a quarter of an inch, the pixel spacing, if you wish.

[Transparency 63]

That was not a surprise to us. We saw these kinds of things before and we did not know what the structure was. There is what happens if you make the resolution 1/8 inch, so now the spatial resolution is half what it was before and now we are beginning to see what looks like scattering from this target below the surface. Of course, the other stuff is the induced particle velocity due to the acoustic wave that is going across the surface.

We have a trigger, we have a phase control here so we can make all these measurements. It takes a very long period of time to make the 128 by 128 measurements with this device; it is rather slow. It does have a set of mirrors that allow you to steer it, but we do have to integrate every one of these, collect every FFT and do the integral.

One of our most serious problems is speckle noise associated with bright and dark spots on the photodetectors and the phase-locked loop detector not working very well.

DR. HARGROVE: What is the integration time?

DR. SABATIER: We integrate from 4 to 5 kHz and that is very fast. The integration time is not important, just a lot of it.

MR. WAKELAND: You said it takes a long time to make the measurements.

DR. SABATIER: I apologize. The reason it takes a long time --

MR. WAKELAND: What is a long time, an hour, a day?

DR. SABATIER: That could take a day. The reason it takes a long time is that we have a lot of speckle noise, where we have bright and dark spots, or dark spots on the photodetectors and when you go through the phase-locked loop detector you typically, to demodulate the signal, you filter it, you give it lots of gain, you make it in a square wave, and when you get poor signal to noise, phase-locked loop detectors do not work.

We have to carefully focus the instrument to make sure that we have a bright spot on the photodetectors, and sometimes we have to move the beam around a little bit or let the wind blow, the wind blows a little bit. The long wavelength is the sound that is leaving the loudspeaker and propagating across the surface and inducing soil particle motion.

The other one -- this is a band of noise that was propagated -- is sound that is backscattered from the target, we think. That is what I think it is. I do not know what it is but I am going to suggest it is sound backscatter from the target.

The velocity amplitudes are small. This does not say what it is. Later, hopefully, one will say what it is. I am reluctant to tell you right now; it may be just volts.

Why are there different wavelengths? That is a good one to explain. When we choose to reconstruct the instantaneous velocity to use the phase, we choose the center of this band. That determines the phase that we use to reconstruct the signal, the velocity. This is RMS but we actually construct the instantaneous velocity.

MR. CURRA: If you backscatter, shouldn't the frequency be the same?

DR. SABATIER: No, we are broadcasting the band of sound and it is backscattering at a preferential frequency or a characteristic frequency.

[Transparency 64]

That was a cylindrical disk 4 inches in diameter, 2 cm below the surface. Here is a triangular disk. We change the spacing, it is quarter-inch resolution, so we are back to the original sized spacing. We chose to lower the frequency because you get a different effect. You do not see the scattering any more. We seem to see something that is probably related to pressure doubling or due to long wavelength scattering from this disk. We were at 5 kHz previously. Now this is at 1 kHz. Same size target.

[Transparency 65]

This is a square target buried below the surface with quarter-inch resolution buried 2 cm below the surface at the same frequency as the triangular target. If you do this with other shapes, that is what you get.

MR. SMITH: A lot of this is like land mines and I do not remember, is that 2-cm depth an average burial depth for those devices?

DR. SABATIER: Some land mines are buried a meter below the surface of the ground. Most of them are all within an inch. We have been able to show you that we can make measurements with the probe microphone down to 20 cm in sands, in beach sands. When you bury land mines, you disturb the soil and it stays disturbed for a long period of time.

We have data that show with geophones we can see them down to 8 cm below the surface for long periods of time. Other people have the same thing, other labs have made these kinds of measurements. We taught other labs how to make these measurements with geophones and they specialize in collecting data. Bury a target, go out and measure it every day for a year. They generate reams of this kind of data and they conclude that you can see them for long periods of time.

The reason is that it takes the soil a long time to go back. We may not be able to tell the difference between a hole that we dug in the ground with a land mine and without a land mine, but this is only the beginning. What we now see is that we probably will be able to see the effects of the structure of the land mine. My intuitive concept is this disk is not ringing.

A land mine has cavities in it, all land mines have cavities in them. In fact, maybe Bonzo can help us with this, or Mark. If you push a land mine in, in the soil, it goes in much easier than it comes out, so it has a spring constant for +X that is different from the spring constant for -X.

So we have interest in doing, say, 800 Hz and 1000 Hz and mixing those 2 and looking at 100 Hz that might come out from a land mine.

You will not see this for a rock, for example, because a rock does not have these properties. The rock does not collapse easier going in than the other way, because it does not have this soft property to it. This has already been demonstrated. People have shown that you can find land mines and you can detect land mines due to the nonlinear property when this object is buried in the ground.

MR. SMITH: An approaching shock in response?

DR. SABATIER: From what little bit I understand, the nonlinear Hooke's law, or the asymmetric, how about that? This is asymmetric flow about the equilibrium.

They all have air spaces in them and they all have leaks, so there may be some Helmholtz response to looking for land mines. Before, we always had the geophone to use to make these measurements. If you push a geophone above 1000 Hz, you have bought the most expensive geophone on the market that I know about. It is easier to buy geophones that work to 80 Hz. It is a lot harder to find geophones that work at very high frequencies and it is hard to get them small enough -- we are seeing things now that we never could see before because we were averaging over a larger spatial area.

We have been funded to take this to a land mine field, we really have. That is what I said I would do to get the funding for this project.

MR. OELZE: How far away would this work? If you are in a land mine field, you do not want to get right up next to the mine to place your speaker and all that.

DR. SABATIER: We have made this laser work at 7 m. The land mine where we are going to go is a land mine without detonators, without trigger mechanisms. (Laughter)

They have the explosives in them. There has been a big effort to make surrogate land mines for radar cross-section, so an immense amount of effort has been spent in trying to construct surrogate land mines that have the appropriate dielectric constants for radar, for ground-penetrating radar.

There are no surrogate land mines that have the correct acoustic properties. Nobody even knows what they should look like yet. What should the signature of a real land mine look like? That is not known. If you take land mines and you ask them to send them to you hollow and you bury them, they really show up, because they are hollow cavities.

MR. YOUNGHOUSE: I am sure this depends on the land mine, but how difficult is it to actually detonate one of those things? If you are trying to detect it using sound waves, is there a chance that the sound waves could cause the thing to explode?

DR. SABATIER: No, we are talking about strain rates of probably of 10^{-8} , or something. Hank Bass says he is going to detonate these with focused sound beams.

There are about 50 million land mines buried on the earth, on the surface of the ground, not the ocean, and we are burying about 30,000 more a year than we find. We bury these -- we do not bury these -- we do. The United States makes very expensive land mines. Our land mines have geophones and microphones on them and processors and they identify the vehicle first and decide if it is friend or foe and then they launch something that is I.R. that goes out and gets anyone. They probably cost a million dollars apiece.

Everybody likes to make little land mines this big that have just enough oomph to take off your leg or your arm and we give these to Mongolia and people fight a war with these things and they plant them everywhere and when they get done they have all these land mines around. And the way they make their living? With agriculture. It is sad, actually.

That is all I have to say

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE 15 Dec 99	3. REPORT TYPE AND DATES COVERED Final 01 Feb 97 - 21 June 98		
4. TITLE AND SUBTITLE Proceedings of the 1998 Physical Acoustics Summer School: Volume I. Transcripts		5. FUNDING NUMBERS PE 61153N G N00014-98-1-0044		
6. AUTHOR(S) Henry E. Bass				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Jamie L. Whitten National Center for Physical Acoustics The University of Mississippi University, MS 38677		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research ONR 331 800 North Quincy Street Arlington, VA 22217-5660		10. SPONSORING / MONITORING AGENCY REPORT NUMBER		
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; Distribution unlimited		12b. DISTRIBUTION CODE		
13. ABSTRACT (Maximum 200 words) Volume I of the Proceedings of the 1998 Physical Acoustics Summer School contains verbatim transcriptions of the lectures given at the summer school. Volume II of these proceedings contains copies of the transparencies used by the lecturers and Volume III contains background materials that were sent to student and discussion leader participants prior to the summer school.				
14. SUBJECT TERMS Physical Acoustics, Acoustic Resonators, Chaos, Nonlinear Bubble Dynamics, Nonlinear Acoustics, Periodic, Random, and Quasiperiodic Media, Porous Media, Quantum Mechanics, Resonant Ultrasound Spectroscopy, Scanning Acoustic Microscopy, Sensor Physics, Sonoluminescence, Thermoacoustics			15. NUMBER OF PAGES 357	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT	

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-1
298-102